

Dell PowerEdge R760

Technical Guide

Notes, cautions, and warnings

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

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

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

Chapter 1: System overview.....	5
Key workloads.....	5
New technologies.....	5
Chapter 2: System features and generational comparison.....	7
Chapter 3: Chassis views and features.....	10
Front view of the system.....	10
Rear view of the system.....	11
Inside the system.....	13
Chapter 4: Processor.....	16
Processor features	16
Supported processors.....	16
Chapter 5: Memory subsystem.....	19
Supported memory.....	19
Chapter 6: Storage.....	20
Storage controllers.....	20
Supported Drives.....	21
Internal storage configuration.....	21
External Storage.....	22
Chapter 7: Networking.....	23
Overview.....	23
OCP 3.0 support.....	23
Supported OCP cards.....	23
OCP NIC 3.0 vs. rack Network Daughter Card comparisons.....	24
Chapter 8: PCIe subsystem.....	25
PCIe risers.....	25
Chapter 9: Power, thermal, and acoustics.....	35
Power.....	35
Power Supply Units.....	36
Thermal.....	38
Thermal design.....	38
Acoustics.....	39
Acoustical configurations of R760	39
PowerEdge acoustical specifications.....	41
Chapter 10: Rack, rails, and cable management.....	49

Rails and cable management information.....	49
Chapter 11: Operating Systems and Virtualization.....	58
Supported Operating Systems.....	58
Chapter 12: Dell Systems Management.....	59
Integrated Dell Remote Access Controller (iDRAC).....	59
Systems Management software support matrix.....	60
Chapter 13: Appendix D: Service and support.....	62
Why attach service contracts.....	62
ProSupport Infrastructure Suite	62
Specialty Support Services.....	64
ProDeploy Infrastructure Suite.....	65
Supplemental Deployment Services	68
Unique Deployment Scenarios.....	69
DAY 2 – Automation Services with Ansible.....	70
Dell Technologies Consulting Services.....	71
Chapter 14: Appendix A: Additional specifications.....	73
Chassis dimension.....	73
Chassis weight.....	74
NIC port specifications.....	74
Video specifications.....	74
USB Ports.....	75
PSU rating.....	76
Environmental specifications.....	77
Thermal restriction matrix.....	79
Thermal air restrictions.....	95
Chapter 15: Appendix A. Standards compliance.....	97
Chapter 16: Appendix C Additional resources.....	98
Appendix A: Dell PowerEdge R760 NEBS level 3.....	99

System overview

The Dell PowerEdge R760 is Dell's latest two-socket, rack server that is designed to run complex workloads using highly scalable memory, I/O, and network options.

The system features:

- Up to 2 x 4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors with up to 56 cores
- Up to 2 x 5th Gen Intel® Xeon® Scalable Processors with up to 64 cores
- Optional Direct Liquid Cooling for required CPU SKU and/or configurations
- 32 DDR5 DIMM slots
- Two redundant AC or DC power supply units
- Up to 12 x 3.5-inch SAS/SATA, or 24 x 2.5-inch, 16 x 2.5-inch, 8 x 2.5-inch, or 2 x 2.5-inch (rear), 4 x 2.5-inch (rear), 4 x EDSFF E3.S (rear) SAS, SATA, or NVMe (HDD/SSD) drives
- Up to 16 x EDSFF E3.S Gen5 NVMe (SSD) drives
- PCI Express® (PCIe) 5.0 enabled expansion slots
- Network interface technologies to cover Network Interface Card (NIC)

Topics:

- [Key workloads](#)
- [New technologies](#)

Key workloads

The Dell PowerEdge R760 offers powerful performance in a purpose-built, cyber resilient, mainstream server. Ideal for:

- Mixed Workload Standardization
- Database and Analytics
- Virtual Desktop Infrastructure
- Artificial Intelligence and Machine Learning

New technologies

Table 1. New technologies

Technology	Detailed Description
5 th Gen Intel® Xeon® Scalable Processors	Core count: Up to 64 core processor
	UPI speed: Up to 4 links per CPU, speed: 12.8 GT/s, 14.4 GT/s, 16 GT/s, 20 GT/s
	Maximum number of PCIe lanes per CPU: Integrated 80 PCIe 5.0 lanes @ 32 GT PCIe Gen5
	Maximum TDP: 350 W
5600 MT/s DDR5 Memory	Max 16 DIMMs per processor and 32 DIMMs per system
	Supports DDR5 ECC RDIMM
4 th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors	Core count: Up to 56 core processor
	UPI speed: Up to 4 links per CPU, speed: 12.8 GT/s, 14.4 GT/s, 16 GT/s
	Maximum number of PCIe lanes per CPU: Integrated 80 PCIe 5.0 lanes @ 32 GT PCIe Gen5

Table 1. New technologies (continued)

Technology	Detailed Description
	Maximum TDP: 350 W
4800 MT/s DDR5 Memory	Max 16 DIMMs per processor and 32 DIMMs per system Supports DDR5 ECC RDIMM
Flex I/O	Optional LOM board, 2x1Gb with BCM5720 LAN controller Rear I/O with: • 1x Dedicated iDRAC Ethernet port • 1 x USB 3.0 • 1 x USB 2.0 • 1 x VGA port (optional for Direct Liquid Cooling configuration) Serial Port Option with STD RIO board Optional OCP Mezz 3.0 (supported by x8 PCIe lanes) Front I/O with: • 1 x USB 2.0 • 1x iDRAC Direct (Micro-AB USB) port • 1 x VGA port
CPLD 1-wire	Support payload data of Front PERC, Riser, BP, and Rear IO to BOSS-N1 and iDRAC
Dedicated PERC	Front Storage module PERC with Front PERC11 & PERC12
Software RAID	OS RAID/S160
Power Supplies	60 mm dimension is the new PSU form factor on a 15G design Titanium 700 W mixed mode HLAC Platinum 800 W mixed mode Titanium 1100 W mixed mode Platinum 1400 W mixed mode Titanium 1400 W mixed mode 277 Vac and HVDC 1100 W -48 V DC Titanium 1800 W mixed mode HLAC 86 mm dimension PSU Platinum 2400 W mixed mode Titanium 2800 W mixed mode HLAC Titanium 3200 W mixed mode 277 Vac and HVDC

System features and generational comparison

The following table shows the comparison between the PowerEdge R760 with the PowerEdge R750.

Table 2. Features comparison

Features	PowerEdge R760	PowerEdge R750
Processors	2 x 4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors 2 x 5th Gen Intel® Xeon® Scalable Processors	2 x 3 rd Generation Intel® Xeon® Processor Scalable Family
CPU interconnect	Intel Ultra Path Interconnect (UPI)	Intel Ultra Path Interconnect (UPI)
Memory	32 x DDR5 RDIMM - Up to 4800 MT/s (1 DPC) / 4400 MT/s (2 DPC) - Up to 5600 MT/s (1DPC) / 4400 MT/s (2 DPC) *	32 x DDR4 RDIMM, LRDIMM 16 x PMem (Intel Optane Persistent Memory 200 Series)
Storage Controllers	- PERC 11G: H755, H755N, H355 - PERC 12G: H965i, H965e - HBA 11: HBA355i, HBA355e - HBA 12: HBA465i, - BOSS-N1 - Software RAID: S160	- PERC 10G: H345, H745, H840 - PERC 11G: H755, H755N, H355 - HBA 11: HBA355i, HBA355e - BOSS-S1 adapter - BOSS-S2 - Software RAID: S150
Drive Bays	Front bays: 3.5 inches, 2.5 inches - 24 Gb SAS, 6 Gb SATA 2.5 inches - Gen3/4 NVMe - EDSFF E3.S - Gen5 NVMe Rear bay: 2.5 inches - 24 Gb SAS, 6 Gb SATA, Gen3/4 NVMe - EDSFF E3.S - Gen5 NVMe	Front bays: 3.5 inches, 2.5 inches - 12 Gb SAS, 6 Gb SATA 2.5 inches - Gen3/4 NVMe Rear bay: 2.5 inches - 12 Gb SAS, 6 Gb SATA, Gen3/4 NVMe
Power Supplies	- AC (Platinum): 800 W, 1400 W, 2400 W - AC (Titanium): 700 W, 1100 W, 1400 W, 1800 W, 2800 W, 3200 W - LVDC @-48 VDC Input: 1100 W	- AC (Platinum): 800 W, 1400 W, 2400 W - AC (Titanium): 700 W, 1100 W - LVDC @-48 VDC Input: 1100 W
Cooling Options	- Air Cooling - Optional Direct Liquid Cooling (DLC)	- Air Cooling - Optional Direct Liquid Cooling (DLC)
	NOTE: DLC is a rack solution and requires rack manifolds and a cooling distribution unit (CDU) to operate.	NOTE: DLC is a rack solution and requires rack manifolds and a cooling distribution unit (CDU) to operate.
Fans	Standard (STD) fans /High performance Silver (HPR Silver) fans/ High performance Gold (HPR Gold) fans	Standard (STD) fans /High performance Silver (HPR Silver) fans/ High performance Gold (HPR Gold) fans
	Up to six hot swap fans	Up to six hot swap fans

Table 2. Features comparison (continued)

Features	PowerEdge R760	PowerEdge R750
Dimension	Height: 86.8 mm (3.41 inches)	Height: 86.8 mm (3.41 inches)
	Width: 482 mm (18.97 inches)	Width: 482 mm (18.97 inches)
	Depth: 772.13 mm (30.39 inches) with bezel	Depth: 772.13 mm (30.39 inches) with bezel
	758.29 mm (29.85 inches) without bezel	758.29 mm (29.85 inches) without bezel
Form Factor	2U rack server	2U rack server
Embedded Management	• iDRAC9 • iDRAC Direct • iDRAC RESTful API with Redfish • iDRAC Service Module • Quick Sync 2 wireless module	• iDRAC9 • iDRAC Direct • iDRAC Service Module • Quick Sync 2 wireless module
Bezel	Optional LCD bezel or security bezel	Optional LCD bezel or security bezel
OpenManage Software	• OpenManage Enterprise • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin • CloudIQ for PowerEdge plug in • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center • OpenManage Integration with Windows Admin Center	• OpenManage Enterprise • OpenManage Power Manager plugin • OpenManage SupportAssist plugin • OpenManage Update Manager plugin
Mobility	OpenManage Mobile	OpenManage Mobile
Integrations	• BMC Truesight • Microsoft System Center • OpenManage Integration with ServiceNow • Red Hat Ansible Modules • Terraform Providers • VMware vCenter and vRealize Operations Manager	• BMC TrueSight • Microsoft System Center • Red Hat Ansible Modules • VMware vCenter
Connections	• IBM Tivoli Netcool/OMNibus • IBM Tivoli Network Manager IP Edition • Micro Focus Operations Manager • Nagios Core • Nagios XI	• IBM Tivoli Netcool/OMNibus • IBM Tivoli Network Manager IP Edition • Micro Focus Operations Manager • Nagios Core • Nagios XI
Security	• Cryptographically signed firmware • Data at Rest Encryption (SEDs with local or external key mgmt) • Secure Boot • Secured Component Verification (Hardware integrity check) • 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	• 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)	2 x 1 GbE LOM
Networking Options	OCP x16 (optional) Mezz 3.0	OCP x8 Mezz 3.0

Table 2. Features comparison (continued)

Features	PowerEdge R760		PowerEdge R750	
	 NOTE: The system allows either LOM card or an OCP card or both to be installed in the system.			
GPU Options	Up to two double wide 350 W, or six single wide 75 W accelerators		Up to two double wide 300 W, or six 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 (optional for Direct Liquid Cooling configuration)	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 (optional for Direct Liquid Cooling configuration)
	Internal Port: 1 x USB 3.0 (optional)		Internal Port: 1 x USB 3.0 (optional)	
PCIe	Up to 8 x PCIe Gen4 or up to 4 x PCIe Gen5 slots		UP to 8 x PCIe Gen4 slots	
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:** * Applicable for 5th Gen Intel® Xeon® Scalable Processors.

Chassis views and features

Topics:

- Front view of the system
- Rear view of the system
- Inside the system

Front view of the system



Figure 1. Front view of 24 x 2.5-inch drive system



Figure 2. Front view of 16 x 2.5-inch drive system (Smart Flow)




Figure 3. Front view of 8 x 2.5-inch drive system



Figure 4. Front view of 12 x 3.5-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 Raid drive system

Rear view of the system



Figure 7. Rear view of the system



Figure 8. Rear view of the system with optional liquid cooling



Figure 9. Rear view of the system with 2 x 2.5-inch rear drive module



Figure 10. Rear view of the system with 4 x 2.5-inch rear drive module



Figure 11. Rear view of the system with 4 x EDSFF E3.S rear drive module

Inside the system




Figure 12. Inside the system




Figure 13. Inside the system with full length risers and GPU shroud




Figure 14. Inside the system with processor liquid cooling module

NOTE: Air shroud is hidden in the above image to show the processor liquid cooling configuration.

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 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
- New Xeon Max processor with integrated 64 GB High Bandwidth Memory (HBM) to increase performance in memory-bound applications

Xeon Max processor modes

1. Xeon Max only mode: Provides best performance when workloads fit into 1 GB/core of capacity and no software changes or DDR required.
2. Flat mode: DDR can be added for workloads needing capacity >1 GB/core with Xeon Max processor and DDR exposed as separate regions and software updates that are needed to optimize performance (higher performance than cache mode).
3. Cache mode: Provides improved performance when workloads need >1 GB/core of capacity with no change to software required and Xeon Max processor caches DDR (symmetric DDR population required).

The 5th Generation Intel® Xeon® Scalable Processor stack is the next generation data center processor offering improved performance, standard increased memory speeds, expanded UPI speeds and enhanced security.

The following lists the features and functions that are in the 5th Generation Intel® Xeon® Scalable Processor offering:

- Increased core counts with up to 64 cores
- Enhanced memory performance with DDR5 and memory speed up to 5600 MT/s in one DIMM per channel (1DPC) and up to 4400 MT/s in two DIMM per channel (2DPC), 24 Gb and 16 Gb DRAM
- Faster UPI with up to four Intel Ultra Path Interconnect (Intel® UPI) at up to 20 GT/s, increasing multi-socket bandwidth
- Enhanced security for virtualized environments with Intel Trust Domain Extensions (Intel® TDX) for confidential computing

Supported processors

The following table shows the Intel Sapphire Rapids (4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors) and Intel Emerald Rapids (5th Gen Intel® Xeon® Scalable Processors) SKUs that are supported on the R760.

Table 3. 4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors supported in R760

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
9480 ¹	1.9	113	16	56	112	Turbo	4800	64 GB	350 W

Table 3. 4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors supported in R760 (continued)

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
9470 ¹	2	105	16	52	104	Turbo	4800	64 GB	350 W
9460 ¹	2.2	98	16	40	80	Turbo	4800	64 GB	350 W
9462 ¹	2.7	75	16	32	64	Turbo	4800	64 GB	350 W
8480+ ¹	2	105	16	56	112	Turbo	4800	4 TB	350 W
8471N ¹	1.8	98	16	52	104	Turbo	4800	4 TB	300 W
8470Q ¹	2.1	105	16	52	104	Turbo	4800	4 TB	350 W
8470N ¹	1.7	98	16	52	104	Turbo	4800	4 TB	300 W
8470 ¹	2	105	16	52	104	Turbo	4800	4 TB	350 W
8468 ¹	2.1	105	16	48	96	Turbo	4800	4 TB	350 W
8460Y+ ¹	2	105	16	40	80	Turbo	4800	4 TB	300 W
8452Y ¹	2	68	16	36	72	Turbo	4800	4 TB	300 W
6454S ¹	2.2	60	16	32	64	Turbo	4800	4 TB	270 W
6430 ¹	2.1	60	16	32	64	Turbo	4800	4 TB	270 W
6414U ¹	2	60	16	32	64	Turbo	4800	4 TB	250 W
8462Y+ ¹	2.8	60	16	32	64	Turbo	4800	4 TB	300 W
6458Q ¹	3.1	60	16	32	64	Turbo	4800	4 TB	350 W
6448Y ²	2.2	60	16	32	64	Turbo	4800	4 TB	225 W
6444Y ¹	3.5	45	16	16	32	Turbo	4800	4 TB	270 W
6442Y ²	2.6	60	16	24	48	Turbo	4800	4 TB	225 W
6438Y+ ²	2	60	16	32	64	Turbo	4800	4 TB	205 W
6438N ²	2	60	16	32	64	Turbo	4800	4 TB	205 W
6438M ²	2.2	60	16	32	64	Turbo	4800	4 TB	205 W
6434 ²	3.7	23	16	8	16	Turbo	4800	4 TB	205 W
6428N ²	1.8	60	16	32	64	Turbo	4800	4 TB	185 W
6426Y ²	2.6	38	16	16	32	Turbo	4800	4 TB	185 W
6421N ²	1.8	60	16	32	64	Turbo	4800	4 TB	185 W
5420+ ²	2	53	16	28	56	Turbo	4400	4 TB	205 W
5418Y ²	2	45	16	24	48	Turbo	4400	4 TB	185 W
5418N ²	1.8	45	16	24	48	Turbo	4400	4 TB	165 W
5416S ²	2	30	16	16	32	Turbo	4400	4 TB	150 W
5415+ ²	2.9	23	16	8	16	Turbo	4400	4 TB	150 W
5412U ²	2.1	45	16	24	48	Turbo	4400	4 TB	185 W
5411N ²	1.9	45	16	24	48	Turbo	4400	4 TB	165 W
4416+ ²	2	38	16	20	40	Turbo	4000	4 TB	165 W
4410Y ²	2	30	16	12	24	Turbo	4000	4 TB	150 W
3408U ²	1.8	23	16	8	16	No Turbo	4000	4 TB	125 W

NOTE: 9480, 9470, 8470Q and 6458Q are supported only in liquid cooling configuration.

Table 4. 5th Gen Intel® Xeon® Scalable Processors supported in R760

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
8592+ ¹	1.9	320	20	64	128	Turbo	5600	4 TB	350 W
8580 ¹	2.0	300	20	60	120	Turbo	5600	4 TB	350 W
8568Y+ ¹	2.3	300	20	48	96	Turbo	5600	4 TB	350 W
8562Y+ ¹	2.8	60	20	32	64	Turbo	5600	4 TB	300 W
8558U ¹	2.0	260	N/A	48	96	Turbo	4800	4 TB	300 W
6548N ¹	2.8	60	20	32	64	Turbo	5200	4 TB	250 W
6548Y+ ¹	2.5	60	20	32	64	Turbo	5200	4 TB	250 W
6542Y ¹	2.9	60	20	24	48	Turbo	5200	4 TB	250 W
6534 ²	3.9	22.5	20	8	16	Turbo	4800	4 TB	195 W
6526Y ²	2.8	37.5	20	16	32	Turbo	5200	4 TB	195 W
5512U ²	2.1	52.5	N/A	28	56	Turbo	4800	4 TB	185 W
4514Y ²	2.0	30	16	16	32	Turbo	4400	4 TB	150 W
4510 ²	2.4	30	16	12	24	Turbo	4400	4 TB	150 W
4509Y ²	2.6	23	16	8	16	Turbo	4400	4 TB	125 W

NOTE: The platform supports Maximum (MAX) and Mainstream (MS) system boards.

- ¹ supports MAX system board
- ² supports MS system board

For more information, see [System board jumpers and connectors](#) section.

Memory subsystem

Topics:

- Supported memory

Supported memory

Table 5. Memory technology comparison

Feature	PowerEdge R760 (DDR5)
DIMM type	RDIMM
Transfer speed	4800 MT/s (1DPC), 4400 MT/s (2DPC) 5600 MT/s (1DPC), 4400 MT/s (2DPC)*
Voltage	1.1 V

NOTE: *Applicable for 5th Gen Intel® Xeon® Scalable Processors.

Table 6. 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
	1 R	16 GB	DDR5 (1.1 V), 5600 MT/s	5600 MT/s	4400 MT/s
	2 R	32 GB, 64 GB, 96 GB	DDR5 (1.1 V), 5600 MT/s	5600 MT/s	4400 MT/s
	4 R	128 GB	DDR5 (1.1 V), 5600 MT/s	5600 MT/s	4400 MT/s
	8 R	256 GB*	DDR5 (1.1 V), 5600 MT/s	5600 MT/s	4400 MT/s

NOTE: 5600 MT/s RDIMMs are applicable for 5th Gen Intel® Xeon® Scalable Processors.

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

NOTE: *256 GB RDIMM with 5th Gen Intel® Xeon® Scalable Processors will be supported in the future release.

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.

NOTE: The size of the RAID 1 drives must be less than that of the second RAID container.

Table 7. PERC Series controller offerings

Performance Level	Controller and Description
Entry	S160
Value	H355, HBA355 (internal/external), HBA465 (internal/external)
Value Performance	H755, H755N
Premium Performance	H965i, H965e Avenger 1 Memory: 8GB DDR4 NV cache 72-bit memory 2133 MHz Low profile form factors Dual A15 1.2 GHz CPU X8PCIe 3.0, x8 12Gb SAS

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 [Storage Controller Manuals](#).

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

NOTE: HBA465e is available post-RTS.

Supported Drives

The table that is shown below lists the internal drives that are supported in R760.

Table 8. 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
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
3.5 inches	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
3.5 inches	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB

Internal storage configuration

R760 available internal storage configurations:

- Zero drive (no backplane)
- 12 x 3.5" (SAS/SATA)
- 12 x 3.5" (SAS/SATA) w/ rear 2 x 2.5" (SAS/SATA)
- 12 x 3.5" (SAS/SATA) w/ rear 2 x 2.5" NVMe Direct
- 12 x 3.5" (SAS/SATA) + 4 x 2.5" (SAS/SATA)
- 12 x 3.5" (SAS/SATA) + 4 x 2.5" (NVMe Direct)
- 12 x 3.5" (SAS/SATA) + 4 x EDSFF E3.S (Gen5 x 4 NVMe Direct)
- 8 x 2.5" NVMe Direct
- 8 x 2.5" (NVMe RAID)
- 8 x 2.5" Universal (SAS/SATA HWRAID + NVMe Direct)
- 16 x 2.5" (NVMe RAID) Smart Flow
- 16 x 2.5" (NVMe Direct) Smart Flow
- 16 x 2.5" (SAS4/SATA) Smart Flow
- 16 x 2.5"(SAS4/SATA)+ 8 x 2.5" NVMe Direct
- 24 x 2.5" (SAS4/SATA) with 8 x Universal slots (SAS/SATA HWRAID + NVMe Direct)
- 24 x 2.5" (SAS4/SATA)
- 24 x 2.5" (SAS4/SATA) + 2 x 2.5" (NVMe Direct)
- 24 x 2.5"(SAS4/SATA) + 2 x 2.5" (SAS/SATA)
- 24 x 2.5"(SAS4/SATA) + 4 x 2.5" (SAS/SATA)
- 24 x 2.5" (SAS4/SATA) + 4 x 2.5" (NVMe Direct)
- 24 x 2.5" (SAS4/SATA) - Dual Controller
- 24 x 2.5" + 2 x 2.5" (SAS4/SATA) - Dual Controller
- 24 x 2.5" (SAS4/SATA) with 8 x Universal slots (SAS/SATA HWRAID + NVMe Direct)
- 24 x 2.5" (SAS4/SATA)
- 24 x 2.5" (SAS4/SATA) - Dual Controller
- 24 x 2.5" (SAS4/SATA) with 4 x Universal slots (SAS/SATA HWRAID + NVMe Direct) + 4 x 2.5" (SAS4/SATA)
- 24 x 2.5 inches (SAS4/SATA) + 4 x EDSFF E3.S (Gen5 x 4 NVMe Direct)
- 16 x 2.5" (8 x SAS4/SATA + 8 x NVMe RAID)

- 16 x EDSFF E3.S (Gen5 x 4 NVMe Direct)
- 16 x EDSFF E3.S (NVMe RAID) - Dual Controller
- 24 x 2.5 inches (NVMe Gen5 switched)
- 24 x 2.5 inches (NVMe RAID - Gen5 Switched) - Dual Controller
- 16 x 2.5" (8 x NVMe RAID + 8 x SAS4/SATA) Smart Flow
- 24 x 2.5" (NVMe Gen4 Direct) Passive
- 8 x 2.5" Universal (SAS/SATA HWRAID + NVMe Direct)

NOTE: The Universal backplane (with universal slot supports SAS/SATA/NVMe drives) supports HW RAID for SAS/SATA with direct attach NVMe, and does not support HW RAID for NVMe.

External Storage

The R760 supports the external storage device types that are listed in the table below.

Table 9. 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

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 10. OCP 3.0 feature list

Feature	OCP 3.0
Form factor	SFF
PCIe Gen	Gen4
Max PCIe width	x8, x16 (with OCP cable)
Max number of ports	4
Port type	BT/SFP/SFP+/SFP28/QSFP56
Max port speed	25 GbE, 100 GbE (with OCP cable)
NC-SI	Yes
SNAPI	Yes
WoL	Yes
Power consumption	15 W–35 W

Supported OCP cards

Table 11. Supported OCP cards

Form factor	Vendor	Port type	Port speed	Port count
OCP 3.0	Broadcom	QSFP56	100 GbE	2
	Mellanox	QSFP56	100 GbE	2
	Intel	SFP28	25 GbE	4
	Broadcom	SFP28	25 GbE	4
	Intel	SFP28	25 GbE	2
	Broadcom	SFP28	25 GbE	2

Table 11. Supported OCP cards (continued)

Form factor	Vendor	Port type	Port speed	Port count
	Mellanox	SFP28	25 GbE	2
	Broadcom	BT	10 GbE	4
	Intel	BT	10 GbE	2
	Intel	BT	10 GbE	4
	Broadcom	BT	10 GbE	2
	Broadcom	BT	1 GbE	4
	Intel	BT	1 GbE	4

NOTE: A 100 GbE OCP card of PCIe width x16 can be used by connecting the OCP cable from SL11_CPU1_PB7 to SL13_CPU1_PB7 on the MAX system board.

NOTE: For storage configurations that already use the SL11_CPU1_PB7 or SL13_CPU1_PB7 connector on the Max system board, there is a restriction on supporting OCP cable.

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 12. 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 is 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

PCIe subsystem

Topics:

- [PCIe risers](#)

PCIe risers

Shown below are the riser offerings for the platform.




Figure 15. Riser connector location on system board

- | | |
|------------|------------|
| 1. Riser 1 | 2. Riser 2 |
| 3. Riser 3 | 4. Riser 4 |




Figure 16. Riser 1B

1. Slot 1
2. Slot 2




Figure 17. Riser 1R

1. Slot 1
2. Slot 2




Figure 18. Riser 1R FL

- 1. Slot 1
- 2. Slot 2




Figure 19. Riser 1P

- 1. Slot 2




Figure 20. Riser 1P FL

1. Slot 2




Figure 21. Riser 1Q

1. Slot 1
2. Slot 2




Figure 22. Riser R1 Paddle




Figure 23. Riser 2A

1. Slot 6
2. Slot 3




Figure 24. Riser 3A

1. Slot 5




Figure 25. Riser 3A FL

1. Slot 5




Figure 26. Riser 3B

1. Slot 4
2. Slot 5




Figure 27. Riser 4P

1. Slot 7




Figure 28. Riser 4P - FL

1. Slot 7




Figure 29. Riser 4B

1. Slot 8
2. Slot 7




Figure 30. Riser 4Q

1. Slot 7
2. Slot 8




Figure 31. Riser 4R

1. Slot 7
2. Slot 8




Figure 32. Riser R4 Paddle

Table 13. PCIe Riser Configurations

Config No.	Riser configuration	No. of Processors	PERC type supported	Rear storage possible
0	NO RSR	2	Front PERC	No
1	R1B+R2A+R3B+R4B	2	Front PERC/PERC Adapter	No
2	R1Q+R2A+R3B+R4Q	2	Front PERC/PERC Adapter	No
3-1	R1P+R2A+R3B+R4P (HL)	2	Front PERC/PERC Adapter	No
3-2	R1P+R2A+R3B+R4P (FL)	2	Front PERC/PERC Adapter	No
4-1	R1P+R2A+R3B+R4R (HL)	2	Front PERC/PERC Adapter	No
5-1	R1R+R2A+R3A+R4P (HL)	2	Front PERC/PERC Adapter	No
5-2	R1R+R2A+R3A+R4P (FL)	2	Front PERC/PERC Adapter	No
6	R2A+R4Q	2	Front PERC/PERC Adapter	Yes

Table 13. PCIe Riser Configurations (continued)

Config No.	Riser configuration	No. of Processors	PERC type supported	Rear storage possible
7	R1Q+R2A+R4Q	2	Front PERC/PERC Adapter	Yes
8	R1B+R2A	1	PERC Adapter	No
9	R1Q+R2A+R4R	1	Front PERC	No
10-1	R1P+R2A+R4R (HL)	1	Front PERC	No
10-2	R1P+R2A+R4R (FL)	1	Front PERC	No
11	R1 Paddle + R2A + R3B + R4 Paddle	2	N/A	No
12	R1Q+R2A+R4Q	2	Front PERC/PERC Adapter	Yes

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 14. 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 Dell EIPT .
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 14. Power tools and technologies (continued)

Feature	Description
	- Power distribution units (PDUs) - Uninterruptible power supplies (UPSs) - Energy Smart containment rack enclosures Find additional information at: Power and Cooling .

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 R760.

Table 15. PSU specifications

PSU	Classes	Heat dissipation (maximum) (BTU/hr)	Frequency (Hz)	AC Voltage			DC Voltage			Current (A)
				200—240 V	100—120 V	277 V	240 V	- (48—60) V	336 V	
700 W mixed mode HLAC	Titanium	2625	50/60	700 W	N/A	N/A	N/A	N/A	N/A	4.1
	N/A	2625	N/A	N/A	N/A	N/A	700 W	N/A	N/A	3.4
800 W mixed mode	Platinum	3000	50/60	800 W	800 W	N/A	N/A	N/A	N/A	9.2—4.7
	N/A	3000	N/A	N/A	N/A	N/A	800 W	N/A	N/A	3.8
1100 W mixed mode	Titanium	4100	50/60	1100 W	1050 W	N/A	N/A	N/A	N/A	12—6.3
	N/A	4100	N/A	N/A	N/A	N/A	1100 W	N/A	N/A	5.2
1400 W mixed mode	Platinum	5250	50/60	1400 W	1050 W	N/A	N/A	N/A	N/A	12—8
	N/A	5250	N/A	N/A	N/A	N/A	1400 W	N/A	N/A	6.6
1400 W mixed mode 277 Vac and HVDC	Titanium	5250	50/60	N/A	N/A	1400 W	N/A	N/A	N/A	5.8
		5250	N/A	N/A	N/A	N/A	N/A	N/A	1400 W	5.17
1800 W mixed mode HLAC	Titanium	6750	50/60	1800	N/A	N/A	N/A	N/A	N/A	10
	N/A	6750	N/A	N/A	N/A	N/A	1800 W	N/A	N/A	8.2
2400 W mixed mode	Platinum	9000	50/60	2400 W	1400 W	N/A	N/A	N/A	N/A	16—13.5
	N/A	9000	N/A	N/A	N/A	N/A	2400 W	N/A	N/A	11.2

Table 15. PSU specifications (continued)

PSU	Class	Heat dissipation (maximum) (BTU/hr)	Frequency (Hz)	AC Voltage			DC Voltage			Current (A)
				200—240 V	100—120 V	277 V	240 V	- (48—60) V	336 V	
2800 W mixed mode HLAC	Titanium	10500	50/60	2800 W	N/A	N/A	N/A	N/A	N/A	15.6
	N/A	10500	N/A	N/A	N/A	N/A	2800 W	N/A	N/A	13.6
1100 W -48 V DC	N/A	4265	N/A	N/A	N/A	N/A	1100 W	N/A	N/A	27
3200 W mixed mode 277 Vac and HVDC	Titanium	12000	50/60	N/A	N/A	3200 W	N/A	N/A	N/A	13
		12000	N/A	N/A	N/A	N/A	N/A	N/A	3200 W	11.5

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 33. PSU power cables




Figure 34. APP 2006G1 power cable

Table 16. PSU power cables

Form factor	Output	Power cable
Redundant 60 mm	700 W mixed mode HLAC	C13
	800 W mixed mode	C13
	1100 W mixed mode	C13

Table 16. PSU power cables (continued)

Form factor	Output	Power cable
	1400 W mixed mode	C13
	1400 W mixed mode 277 Vac and HVDC	APP 2006G1
	1800 W mixed mode HLAC	C15
Redundant 86 mm	2400 W mixed mode	C19
	2800 W mixed mode HLAC	C21
	3200 W mixed mode 277 Vac and HVDC	APP 2006G1

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

NOTE: C13 power cable combined with C14 to C15 jumper power cable can be used to adapt a 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 35. Thermal design characteristics

The thermal design of the PowerEdge R760 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. A 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 a 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 R760 Installation and Service Manual at [PowerEdge Manuals](#) and “Advanced Thermal Control: Optimizing across Environments and Power Goals” on Dell.com.
- Cooling redundancy: The R760 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the R760 reliable under a wide range of operating environments.

Acoustics

Acoustical configurations of R760

Dell PowerEdge R760 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 17. Configurations tested for acoustical experience

Configuration	Quietest GPU configuration	Entry/ Quietest configuration	Typical-1, 2.5-inch	Typical-2, 3.5-inch	GPU configuration	NVMe Box
CPU TDP	125 W	125 W	165 W	165 W	205 W	300 W
CPU Quantity	2	2	2	2	2	2
RDIMM Memory	16 GB DDR5	16 GB DDR5	16 GB DDR5	32 GB DDR5	32 GB DDR5	16 GB DDR5
Memory Quantity	8	8	16	16	32	32
Backplane Type	8 x 2.5-inch BP	12x 3.5-inch BP	8 x 2.5-inch exp BP x2	12 x 3.5-inch BP + 2 x 2.5-inch rear BP	16 x 2.5-inch exp BP	24 x 2.5-inch exp BP (NVMe)
HDD Type	X	SATA 3.5-inch 4 TB	X	12 x 3.5-inch front 12 TB, 2 x 2.5-inch rear SSD	X	X
HDD Quantity	X	2	X	12 + 2	X	X
Flash Drives	PCIe SSD	X	PCIe SSD	X	PCIe SSD	PCIe SSD
Flash Quantity	8	X	8	X	16	24
PSU Type	1400 W	800 W	800 W	1400 W	2400 W	2400 W
PSU Quantity	2	2	2	2	2	2
OCP	2x10 G	2x10 G	10/25 2-port	10/25 2-port	10/25 2-port	2x25 G
PCI 1	X	H355	X	H755	X	X
PCI 2	X	X	X	X	GPU	X
PCI 3	X	X	X	X	X	X
PCI 4	X	X	2-port 25 Gb	2-port 10 Gb	2-port 25 Gb	X
PCI 5	X	X	2-port 25 Gb	2-port 10 Gb	2-port 25 Gb	100 Gb PCI

Table 17. Configurations tested for acoustical experience (continued)

Configuration	Quietest GPU configuration	Entry/ Quietest configuration	Typical-1, 2.5-inch	Typical-2, 3.5-inch	GPU configuration	NVMe Box
PCI 6	25/50 Gb	X	X	X	X	X
PCI 7	A30	X	X	X	GPU	100 Gb PCI
PCI 8	X	X	X	X	X	X
PERC	Front H755n	Adapt H355	Front H7455n	Adapt H755	Front H755n	Front H755n

Table 18. Acoustical experience of R760 configurations

Configuration		Quietest GPU configuration	Entry/ Quietest configuration	Typical-1, 2.5-inch	Typical-2, 3.5-inch	GPU configuration	NVMe Box
Acoustical Performance: Idle/ Operating @ 25°C Ambient							
L _{wA,m} (B)	Idle ⁽⁴⁾	6.5	5.1	5.5	6.4	6.9	6.8
	Operating/ Customer usage operating ⁽⁵⁾⁽⁶⁾	8.1	5.1	5.5	6.4	8.5	6.8
K _v (B)	Idle ⁽⁴⁾	0.4	0.4	0.4	0.4	0.4	0.4
	Operating/ Customer usage operating ⁽⁵⁾⁽⁶⁾	0.4	0.4	0.4	0.4	0.4	0.4
L _{pA,m} (dB)	Idle ⁽⁴⁾	51	36	41	48	55	54
	Operating/ Customer usage operating ⁽⁵⁾⁽⁶⁾	69	36	41	48	74	54
Prominent discrete tones ⁽³⁾		Prominence ratio ≤ 17 dB	No audible tones		Prominence ratio < 15 dB	Prominence ratio ≤ 17 dB	Prominence ratio < 15 dB
Acoustical Performance: Idle @ 28°C Ambient							
L _{wA,m} ⁽¹⁾ (B)		7.3	5.4	5.9	6.7	7.3	7.1
K _v (B)		0.4	0.4	0.4	0.4	0.4	0.4
L _{pA,m} ⁽²⁾ (dB)		59	36	45	52	59	57
Acoustical Performance: Max. loading @ 35°C Ambient							
L _{wA,m} ⁽¹⁾ (B)		9.0	6.0	7.0	7.8	9.0	7.8
K _v (B)		0.4	0.4	0.4	0.4	0.4	0.4
L _{pA,m} ⁽²⁾ (dB)		79	44	58	66	79	65

⁽¹⁾L_{wA, m}: The declared mean A-weighted sound power level (L_{wA}) 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 the ISO 7779 declaration requirement.

⁽²⁾L_{pA, 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 the ISO 7779 declaration requirement.

⁽³⁾Prominent tones: Criteria of Annex D of ECMA-74 and the 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 19. 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			N/A

Table 19. 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
		○ 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 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 20. 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 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 21. 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.			

Table 21. 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
		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 22. 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

Table 22. 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	
	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 23. 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 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 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 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

Rack, rails, and cable management

Topics:

- [Rails and cable management information](#)

Rails and cable management information

The rail offerings for the PowerEdge R760 consist of two general types: sliding and static. 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 [Dell site](#) 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 are two types of sliding rails available, ReadyRails II sliding rails and 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).

B21 ReadyRails sliding rails for 4-post racks

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

NOTE: For situations where CMA support is not required, the outer CMA mounting brackets can be uninstalled from the sliding rails. This reduces the overall length of the rails and eliminates the potential interference with rear-mounted PDUs or the rear rack door.



Figure 36. Sliding rails with optional CMA



Figure 37. Sliding rails with optional SRB

B22 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 an optional cable management arm (CMA).
- Support for optional strain relief bar (SRB).

NOTE: For situations where CMA support is not required, the outer CMA mounting brackets can be uninstalled from the sliding rails. This reduces the overall length of the rails and eliminates the potential interference with rear-mounted PDUs or the rear rack door.

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

B20 static rails summary

The static rails offer a greater adjustability range and a smaller overall mounting footprint than the sliding rails because of their reduced complexity and lack of need for CMA support. The static rails support a wider variety of racks than the sliding rails. However, they do not support serviceability in the rack and are thus not compatible with the CMA. The static rails are also not compatible with SRB.



Figure 38. Static rails

Static rails features summary

Static rails for 4-post and 2-post racks:

- Supports Stab-in installation of the chassis to the rails.
- Support tool-less installation in 19" EIA-310-E compliant square or unthreaded round hole 4-post racks including all generations of Dell racks.
- Support tooled installation in 19" EIA-310-E compliant threaded hole 4-post and 2-post racks.
- Support for tooled installation in Dell Titan or Titan-D rack.

NOTE:

- Screws are not included with the static rail kit since racks are offered with various thread types. The screws are provided for mounting static rails in racks with threaded mounting flanges.
- Screw head diameter should be 10 mm or less.

2-Post racks installation

If installing to 2-Post (Telco) racks, the ReadyRails II static rails (B20) must be used. Sliding rails support mounting in 4-post racks only.



Figure 39. Static rails in 2-post center mount configuration

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 (B22) 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. For tooled installation, Stab-in Static rails (B20) must be used for bezel alignment with storage systems.

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.

NOTE: CMA is not supported in the Direct Liquid Cooling configuration.

The CMA can be mounted to either side of the sliding rails without the use of tools or the need for conversion. 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 40. Sliding rails with CMA




Figure 41. CMA Cabling

Strain Relief Bar (SRB)

The optional strain relief bar (SRB) for the PowerEdge R760 organizes and supports cable connections at the rear end of the server to avoid damage from bending.




Figure 42. 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 the system into the rack (option A: Drop-In)

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



Figure 43. 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 44. 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 45. 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 46. Pull out the intermediate rail

Table 24. Rail component label

Number	Component
1	Intermediate rail

Table 24. Rail component label (continued)

Number	Component
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 47. Attach the inner rails to the system

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



Figure 48. 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 49. Slide system into the rack

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](#).

Dell Systems Management

Dell delivers management solutions that help IT administrators effectively deploy, update, monitor, and manage IT assets. Dell 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](#).

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 25. 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 26. 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 26. 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 2019 or 2022, Ubuntu, CentOS	Supported (Tier-1)

Appendix D: Service and support

Topics:

- [Why attach service contracts](#)
- [ProSupport Infrastructure Suite](#)
- [Specialty Support Services](#)
- [ProDeploy Infrastructure Suite](#)
- [Supplemental Deployment Services](#)
- [Unique Deployment Scenarios](#)
- [DAY 2 – Automation Services with Ansible](#)
- [Dell Technologies Consulting Services](#)

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 failure rates for servers are roughly 1% and more commonly, customers seek 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 (up to 12 years: including seven years standard support and an additional five years of Post-Standard Support). Details of the ProSupport Suite and benefits are listed below.

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.

ProSupport Infrastructure Suite | Enhanced value across all offers!

	Basic Hardware Support	ProSupport for Infrastructure	ProSupport Plus for Infrastructure	Changes with August 2023 release
Technical support availability and response objective	9/5, immediate	24/7, immediate	24/7, immediate	No change
Covered products	Hardware	Hardware & Software	Hardware & Software	No change
Onsite response service level	NBD	NBD or 4-hour	4-hour	ProSupport Plus NBD is retired
ProSupport AIOps platforms	•	•	•	MyService360 and TechDirect (all offers) CloudIQ (ProSupport & ProSupport Plus)
Dell Security Advisories	•	•	•	Available on additional products
Proactive issue detection with automated case creation	•	•	•	New to Basic
Predictive hardware anomaly detection		•	•	New to ProSupport
Access to software updates		•	•	No change
CloudIQ health and cybersecurity monitoring & analytics		•	•	Enhanced features
Incident Manager for Severity 1 cases		•	•	No change
Mission Critical support			•	Enhanced features
Priority access to remote senior support engineers ¹			•	No change
Service Account Manager			•	No change
Proactive system maintenance			•	No change
Limited 3 rd party software support ²			•	No change

¹Based on availability

²Software license can be purchased through Dell or BYOL - see Service Descriptions for details.

Dell Technologies

Figure 50. ProSupport Enterprise Suite

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. Service Account 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

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

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 CSPs (Cloud Serviced Providers)**

ProSupport One for CSPs is a unique offer that is designed for a limited set of Dell accounts purchasing Gen AI computing solutions greater than 1,000 servers and \$250M in sales. PS1 for CSPs improves the entire services experience combining support, deployment (rack integration), residency services, a designated support engineer and the LOIS parts locker as one holistic bundle. Special pricing has been determined to compete effectively against competitors and provide the best customer experience. PS1 for CSPs 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 special bundle offer. More details on PS1 for CSPs [here](#).

- **Logistics Online Inventory Solution (LOIS)**

Ideal for large organizations that have their own staff to support their data center. Dell offers a service that is called Logistics Online Inventory Solution which is an onsite parts locker that provides self-maintainers with a local inventory of common replacement components. Having access to these parts lockers allows the self-maintainer to replace a failed component immediately without delay. 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 5 offers: ProDeploy Configuration Services, ProDeploy Rack Integration Services, Basic Deployment, ProDeploy, and ProDeploy Plus.

ProDeploy Infrastructure Suite

Versatile choices for accelerated deployments




Figure 51. ProDeploy Infrastructure Suite

Factory-based Services

The new Factory Services consist of two tiers of deployment that happen prior to shipping to the Customer's site.

ProDeploy FLEX 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.

ProDeploy FLEX Rack Integration

Ideal for customers seeking to build out fully integrated racks prior to shipping. These rack builds include hardware install, cabling, and full system configuration. You can also add-on a factory stress test and optional on-site final rack configuration to complete the rack installation.

- STANDARD SKUs for Rack Integration is available in USA only and requires:
 - 20 or more devices (R and C series servers, VxRail, and all Dell or non-Dell switches.
 - Shipping to contiguous USA.
- USE CUSTOM QUOTE for Rack Integration scenarios that require:
 - Shipment to any country or region outside USA or shipping outside contiguous USA
 - Shipping to multiple locations
 - Racks containing fewer than 20 servers
 - Any rack that includes Storage.

ProDeploy Flex | Modular deployment (built in factory, onsite or remote)

Pre -deployment	Single point of contact for project management	●
	Expanded end-to-end project management	Selectable
	Site readiness review and implementation planning	●
Deployment	Deployment service hours	24/7
	Hardware installation options ¹	Onsite, factory ^{2,5} or remote ³
	System software installation and configuration options ¹	Onsite, factory ^{2,5} or remote ³
	Multivendor networking deployment ⁴	Onsite, factory ^{2,5} or remote ³
	Onsite Deployment in remote locations	Selectable
	Onsite Deployment in challenging environments	Selectable
	Onsite Deployment with special site-based protocols or requirements	Selectable
	Install connectivity software based on Secure Connect Gateway technology	●
	Dell NativeEdge Orchestrator deployment	Selectable
	Configure 3 rd party software applications and workloads ⁴	Selectable
Post -deployment	Deployment verification, documentation, and knowledge transfer	●
	Configuration data transfer to Dell support	●
Online collaboration	Online collaborative environment - Planning, managing and tracking delivery process	●

¹ Hardware and Software delivery methods can be independently chosen; selecting Rack integration for software requires hardware Rack integration to also be selected.

² Factory Rack Integration for server and VxRail; includes associated Dell network switches; final onsite rack installation available.

³ Remote hardware option includes project specific instructions, documentation and live expert guidance for hardware installation.

⁴ Select 3rd party multivendor networking and software applications.

⁵ Pair with Field Onsite Hardware service for final installation.

Figure 52. ProDeploy Flex modular services

Field-based services

• 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 perform 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 and 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 implementation includes everything that is mentioned in ProDeploy Plus except it does not include the added value, cybersecurity implementation and best practices.

ProDeploy Infrastructure Suite | Field services

		Basic Deployment	ProDeploy	ProDeploy Plus
Pre-deployment	Single point of contact for project management	-	•	In region
	Site readiness review and implementation planning	-	•	•
Deployment	Deployment service hours	Business hours	24/7	24/7
	Hardware installation options	Onsite	Onsite or guided ¹	Onsite
	System software installation and configuration options	-	Remote	Onsite
	Install connectivity software based on Secure Connect Gateway technology ²	-	•	•
	Implement CyberSecurity best practices and policies in APEX AI/Ops Infrastructure Observability	-	-	•
Post-deployment	Deployment verification, documentation and knowledge transfer	-	•	•
	Configuration data transfer to Dell technical support	-	•	•
Online collaboration	Online collaborative platform in TechDirect for planning, managing and tracking delivery	-	•	•

¹ Choose from onsite hardware installation or a guided option including project specific instructions, documentation and live expert guidance
² Post deployment use for intelligent, automated support & insights

Figure 53. 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.

ProDeploy FLEX

ProDeploy Flex is a modular service and a powerful tool for you to attach more services and improve revenue and margins. The ProDeploy Flex modular offer allows sales teams to build and better tailor services by mixing factory and field delivery options. You can also select special deployment scenarios without going to the custom order desk. FLEX is ideal for unique deployments where ProDeploy or ProDeploy Plus are not an adequate answer to the customer needs. Key features of ProDeploy FLEX :

- Build deployment quotes using modular, selectable features for both hardware and software.
- The system automatically scales pricing based on volume.
- Ideal for customers who require NativeEdge Orchestrator or edge deployments.
- Ability to add deployment services to third-party networking devices.

Deployment of HPC

High-Performance Computing (HPC) implementations require specialists that understand advanced feature sets. Dell deploys the world 's fastest systems and understands the nuances that make them perform. HPC deployments are most often scoped as custom service engagements, however we can do smaller HPC clusters under 300 nodes using a standard ProDeploy SKU. Any standard SKU for HPC deployment will be sold as one base SKU per cluster (ProDeploy for HPC Base) along with one ProDeploy for HPC Add-on for each device in the cluster (server nodes and switches).

Scope of ProDeploy for HPC:

 **NOTE:** Available as standard SKUs in US and Canada. Custom Service would be required for all other regions.





Figure 54. Standard deliverables of ProDeploy for HPC

Build HPC solutions for your unique requirements

Choose ProDeploy for HPC or Custom deploy

ProDeploy service includes configuration of most OS, cluster mgmt., networking and benchmarking



Notes related to networking above: Omni-Path is no longer an Intel Product, but is now distributed by a company called Cornelis, and Mellanox was purchased by Nvidia, and now goes by Nvidia Networking.

Figure 55. Visual view of HPC deployment options to include hardware and software

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 portfolio and partner ecosystem of Dell Technologies to help achieve real business outcomes. From multi cloud, 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 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.

Managed	Outsourcing or CAPEX model	APEX	as-a-Service or OPEX model
We manage your technology using our people and tools.¹ • Managed detection and response* • Technology Infrastructure • End-user (PC/desktop) • Service desk operations • Cloud Managed (Pub/Private) • Office365 or Microsoft Endpoint		We own all technology so you can off-load all IT decisions. • APEX Cloud Services • APEX Flex on Demand elastic capacity • APEX Data Center Utility pay-per-use model	

1 – Some minimum device counts may apply. Order via: ClientManagedServices.sales@dell.com

* Managed detection and response covers the security monitoring of laptops, servers, & virtual servers. Min. 50 devices combined. No Networking or Storage-only systems [SAN/NAS]. Available in 32 countries. [Details here](#)

Figure 56. Dell Managed Services

Managed Detection and Response (MDR)

Dell Technologies Managed Detection and Response (MDR) is powered by Secureworks Taegis XDR software platform. MDR is a managed service that secures the customer’s IT environment against malicious actors and provides remediation if and when a threat is identified. When a customer purchases MDR, they will receive the following features from our team:

- Dell badge resources
- Agent rollout assistance to help deploy the Secureworks Endpoint Agent
- 24x7 threat detection & investigation
- Up to 40hrs per quarter of response and active remediation activities
- If the customer experiences a breach, we will provide up to 40hrs per year of Cyber incident response initiation
- Quarterly reviews with the customer to review the data

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 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](https://education.dell.com).

Appendix A: Additional specifications

Topics:

- Chassis dimension
- Chassis weight
- NIC port specifications
- Video specifications
- USB Ports
- PSU rating
- Environmental specifications

Chassis dimension

The R760 has the following dimensions:




Figure 57. Chassis dimensions

Table 27. Chassis dimensions

Model number	Xa	Xb	Y	Za with bezel	Za without bezel	Zb	Zc	Max Sys Wgt	Chassis
R760	482 mm	434 mm	86.8 mm	35.84 mm	22 mm	700.7 mm	736.29 mm	36.1 kg	2U

Chassis weight

Table 28. Chassis weight

System Configuration	Maximum Weight
A server with fully populated drives	36.1 kg (79.58 lbs)
A server without drives and PSU installed	25.1 kg (55.33 lbs)

NIC port specifications

The PowerEdge R760 system supports up to two Network Interface Controller (NIC) ports embedded on the LAN on Motherboard (LOM) card and up to four ports integrated on the Open Compute Project (OCP) NIC card.

Table 29. NIC port specification for the system

Feature	Specifications
LOM card (optional)	1 GbE x 2
OCP NIC card (OCP NIC 3.0) (optional)	1GbE x 4, 10 GbE x 2, 10 GbE x 4, 25 GbE x 2, 25 GbE x 4, 100GbE x 2
Management Interface Card (MIC) to support Dell Data Processing Unit (DPU) card (optional)	25 GbE x 2, 100 GbE x 2 or 200 GbE x 2

- NOTE:** The system allows either LOM card or an OCP NIC card or both to be installed in the system.
- NOTE:** On the MS system board, the supported OCP NIC PCIe width is x8; when x16 PCIe width is installed, it is downgraded to x8.
- NOTE:** A 100 GbE OCP NIC card of PCIe width x16 can be used by connecting the OCP NIC cable from SL11_CPU1_PB7 to SL13_CPU1_PB7 on the MAX system board.
- NOTE:** For storage configurations that already use the SL11_CPU1_PB7 or SL13_CPU1_PB7 connector on the Max system board, there is a restriction on supporting OCP NIC cable.
- NOTE:** The system allows either LOM card or MIC card to be installed in the system.

Video specifications

The PowerEdge R760 system supports integrated Matrox G200 graphics controller with 16 MB of video frame buffer.

Table 30. Video specifications for R760

Resolution	Refresh rate (Hz)	Color depth (bits)
1024 x 768	60	8, 16, 32
1280 x 800	60	8, 16, 32
1280 x 1024	60	8, 16, 32

Table 30. Video specifications for R760 (continued)

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

USB Ports



Figure 58. Front USB Port



Figure 59. Rear USB Port




Figure 60. Internal USB Port

Table 31. Systems USB Specifications

Front		Rear		Internal	
USB port type	No. of ports	USB port type	No. of ports	USB port type	No. of ports
USB x.2.0 – compliant port	1	USB x.2.0 – compliant port	1	USB x.3.0 – compliant port	1
		USB x.3.0 – compliant port	1		

PSU rating

Below table lists the power capacity of the PSUs in high/low line operation mode.

Table 32. PSUs highline and lowline ratings

PSU	700 W Titanium	800 W Platinum	1100 W Titanium	1100 W -48 VDC	1400 W Platinum	1400 W Titanium	1800 W Titanium	2400 W Platinum	2800 W Titanium	3200 W Titanium
Peak Power (Highline /-72 VDC)	1190 W	1360 W	1870 W	1360 W	2380 W	2380 W	3060 W	4080 W	4760 W	5440 W
Highline /-72 VDC	700 W	800 W	1100 W	1100 W	1400 W	1400 W	1800 W	2400 W	2800 W	3200 W
Peak Power (Lowline /-40 VDC)	N/A	1360 W	1785 W	1785 W	1785 W	1785 W	N/A	2380 W	N/A	N/A
Lowline/-40 VDC	N/A	800 W	1050 W	1100 W	1050 W	1050 W	N/A	1400 W	N/A	N/A
Highline 240 VDC	700 W	800 W	1100 W	N/A	1400 W	1400 W	1800 W	2400 W	2800 W	3200 W
DC-(48 —60) V	N/A	N/A	N/A	1100 W	N/A	N/A	N/A	N/A	N/A	N/A

The PowerEdge R760 supports up to two AC power supplies with 1+1 redundancy, autosensing, and auto switching capability.

If two PSUs are present during POST, a comparison is made between the wattage capacities of the PSUs. In case the PSU wattages do not match, the larger of the two PSUs is enabled. Also, there is a PSU mismatch warning that is displayed in the BIOS, iDRAC, or on the system LCD.

If a second PSU is added at run-time, in order for that particular PSU to be enabled, the wattage capacity of the first PSU must equal the second PSU. Otherwise, the PSU is flagged as unmatched in iDRAC and the second PSU is not enabled.

Dell PSUs have achieved Platinum efficiency levels as shown in the table below.

Table 33. PSU efficiency level

Efficiency Targets by Load						
Form factor	Output	Class	10%	20%	50%	100%
Redundant 60 mm	700 W AC	Titanium	90.00%	94.00%	96.00%	91.50%
	800 W AC	Platinum	89.00%	93.00%	94.00%	91.50%

Table 33. PSU efficiency level (continued)

Efficiency Targets by Load						
Form factor	Output	Class	10%	20%	50%	100%
	1100 W AC	Titanium	90.00%	94.00%	96.00%	91.50%
	1100 W -48 V DC	N/A	85.00%	90.00%	92.00%	90.00%
	1400 W AC	Platinum	89.00%	93.00%	94.00%	91.50%
	1400 W AC	Titanium	90.00%	94.00%	96.00%	91.50%
	1800 W AC	Titanium	90.00%	94.00%	96.00%	94.00%
Redundant 86 mm	2400 W AC	Platinum	89.00%	93.00%	94.00%	91.50%
	2800 W AC	Titanium	90.00%	94.00%	96.00%	94%
	3200 W AC	Titanium	90.00%	94.00%	96.00%	91%

Environmental specifications

 **NOTE:** For additional information about environmental certifications, refer to the *Product Environmental Datasheet* located with the *Documentation* on [Dell Support](#).

Table 34. Continuous Operation Specifications for ASHRAE A2

Temperature	Specifications
Allowable continuous operations	
Temperature range for altitudes <= 900 m (<= 2953 ft)	10–35°C (50–95°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) 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 m (1.8°F/984 Ft) above 900 m (2953 Ft)

Table 35. Continuous Operation Specifications for ASHRAE A3

Temperature	Specifications
Allowable continuous operations	
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–40°C (41–104°F) with no direct sunlight on the equipment
	Excursion Limited Operation
	5-35°C (41-95°F) Continuous Operation 35-40°C (95-104°F) 10% Annual Runtime
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 85% RH with 24°C (75.2°F) maximum dew point
Operational altitude de-rating	Maximum temperature is reduced by 1°C/175 m (1.8°F/574 Ft) above 900 m (2953 Ft)

Table 36. Continuous Operation Specifications for ASHRAE A4

Temperature	Specifications
Allowable continuous operations	

Table 36. Continuous Operation Specifications for ASHRAE A4 (continued)

Temperature	Specifications	
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–45°C (41–113°F) with no direct sunlight on the equipment	
	Excursion Limited Operation	5–35°C (41–95°F) Continuous Operation
		35–40°C (95–104°F) 10% Annual Runtime
		40–45°C (104–113°F) 1% Annual Runtime
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point	
Operational altitude de-rating	Maximum temperature is reduced by 1°C/125 m (1.8°F/410 Ft) above 900 m (2953 Ft)	

Table 37. Common Environmental Specifications for ASHRAE A2, A3 and A4

Temperature	Specifications
Allowable continuous operations	
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 <i>i</i> NOTE: * - Per ASHRAE thermal guidelines for tape hardware, these are not instantaneous rates of temperature change.
Non-operational temperature limits	-40 to 65°C (-40 to 149°F)
Non-operational humidity limits	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	3,050 meters (10,006 feet)

Table 38. Maximum vibration specifications

Maximum vibration	Specifications
Operating	0.21 G _{rms} at 5 Hz to 500 Hz for 10 minutes (all operation orientations)
Storage	1.88 G _{rms} at 10 Hz to 500 Hz for 15 minutes (all six sides tested)

Table 39. Maximum shock pulse specifications

Maximum shock pulse	Specifications
Operating	Six consecutively executed shock pulses in the positive and negative x, y, and z axis of 6 G for up to 11 ms
Storage	Six consecutively executed shock pulses in the positive and negative x, y, and z axis (one pulse on each side of the system) of 71 G for up to 2 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 may need to rectify the environmental conditions. Remediation of environmental conditions is the responsibility of the customer.

Table 40. Particulate contamination specifications

Particulate contamination	Specifications
Air filtration	Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit

Table 40. Particulate contamination specifications (continued)

Particulate contamination	Specifications
	NOTE: This condition applies to data center environments only. Air filtration requirements do not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor. 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.
Walk-Up Edge Data Center or Cabinet (sealed, closed loop environment)	Filtration is not required for cabinets that are anticipated to be opened 6 times or less per year. Class 8 per ISO 1466-1 filtration as defined above is required otherwise NOTE: In environments commonly above ISA-71 Class G1 or that may have known challenges, special filters may be required.

Table 41. Gaseous contamination specifications

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

Thermal restriction matrix

Table 42. Processor and heat sink matrix

Heat sink	Processor TDP
STD HSK	≤ 165 W (supports only 2.5-inch drives and non-GPU configuration)
2U HPR HSK	125 W–250 W (supports 3.5-inch drives and non-GPU configuration)
	165 W–350 W (supports 2.5-inch drives and non-GPU configuration)
L-type HSK	Supports all GPU/FPGA configurations

NOTE: All GPU/FPGA cards require 1U L-type HSK and GPU shroud.

Table 43. Label reference

Label	Description
STD	Standard
HPR (Silver)	High performance Silver (HPR Silver) fan
HPR (Gold)	High performance Gold (HPR Gold) fan

Table 43. Label reference (continued)

Label	Description
HSK	Heat sink
LP	Low profile
FH	Full height
DLC	Direct Liquid Cooling

NOTE: The ambient temperature of the configuration is determined by the critical component in that configuration. For example, if the processor's supported ambient temperature is 35°C (95°F), the DIMM is 35°C (95°F), and the GPU is 30°C (86°F), the combined configuration can only support 30°C (86°F).

Thermal restriction matrix for 4th Gen Intel® Xeon® Scalable or Intel® Xeon® Max Processors

Table 44. Thermal restriction matrix for air cooled configuration

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch		Ambient temperature
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan									HPR GOLD fan 70%^	
3408U	125 W ¹	8	79	STD	STD	STD	STD	STD	HPR SLVR	STD	HPR GOLD	HPR SLVR	HPR GOLD	35°C (95°F)
5415+	150 W ¹	8	78	STD	STD	STD	STD	STD	HPR SLVR	STD	HPR GOLD	HPR SLVR	HPR GOLD	35°C (95°F)
4410Y		12	78											
5416S		16	78											
5418N	165 W ¹	24	84	STD	STD	STD	STD	STD	HPR SLVR	STD	HPR GOLD	HPR SLVR	HPR GOLD	35°C (95°F)
5411N		24	84											
4416+		20	82											
6426Y	185 W ¹	16	72	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD	HPR GOLD	35°C (95°F)
5418Y		24	80											
5412U		24	80											
6428N		32	85											
6421N		32	85											

Table 44. Thermal restriction matrix for air cooled configuration (continued)

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch		Ambient temperature
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan									HPR GOLD fan 70% [^]	
6434	205 W ¹	8	96	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD	HPR GOLD	35°C (95°F)
5420+		28	84											
6438Y+		32	76											
6438M		32	84											
6438N		32	84											
6442Y	225 W ¹	24	79	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD*	HPR GOLD*	35°C (95°F)
6448Y		32	79											
6444Y	270 W ²	32	75	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD	Require d DLC	Require d DLC	35°C (95°F)
8462Y+	300 W ²	32	81	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR fan	HPR SLVR	HPR GOLD	Require d DLC	Require d DLC	35°C (95°F)
6458Q	350 W ²	32	64	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	Require d DLC	35°C (95°F)
6414U	250 W ²	32	76	STD fan	STD fan	STD fan	STD fan	STD fan	HPR SLVR fan	HPR SLVR	HPR GOLD	HPR GOLD*	HPR GOLD*	35°C (95°F)
6454S	270 W ²	32	71	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR fan	HPR SLVR	HPR GOLD	Require d DLC	Require d DLC	35°C (95°F)
6430		32	71											
8471N	300 W ²	52	76	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD	Require d DLC	Require d DLC	35°C (95°F)
8470N		52	76											
8460Y+		40	75											
8452Y		36	75											

Table 44. Thermal restriction matrix for air cooled configuration (continued)

Configuration				No back plane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch		Ambient temperature
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan									HPR GOLD fan 70%^	
8480+	350 W ²	56	79	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD*	Required DLC	Required DLC	35°C (95°F)
8470		52	79											
8468		48	79											
8470Q	350 W ²	52	57	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	35°C (95°F)
9480	350 W ²	56	64	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	35°C (95°F)
9470		52	64											
9460	350 W ²	40	77	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD*	Required DLC	Required DLC	35°C (95°F)
9462		32	77											

NOTE: The platform supports Maximum (MAX) and Mainstream (MS) system boards.

- ¹ supports MS system board (CPU TDP < 250 W)
- ² supports MAX system board (CPU TDP => 250 W)

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

NOTE: *Supported ambient temperature is 30°C (86°F).

Table 45. Thermal restriction matrix for memory with air cooled configuration (non-GPU)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch		
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
DIMM Configuration	2DPC/Power	STD fan (CPU TDP <= 250 W)						HPR SLVR fan (CPU TDP up to 350 W)	STD fan (CPU TDP <= 165 W)	HPR GOLD fan (CPU TDP up to 350 W)	HPR SLVR fan 70% (CPU TDP up to 165 W)^	
256 GB RDIMM	12.7 W	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	35°C (95°F)	Required DLC	35°C (95°F)	Required DLC	Required DLC	
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	
DIMM Configuration	2DPC/Power	HPR SLVR fan (CPU TDP up to 350 W)								HPR GOLD fan (CPU TDP up to 350 W)	HPR GOLD fan 70% (CPU TDP up to 250 W)	
256 GB RDIMM	12.7 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	Required DLC	Required DLC
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

Table 45. Thermal restriction matrix for memory with air cooled configuration (non-GPU) (continued)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch	
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan
DIMM Configuration	2DP C/Power	STD fan (CPU TDP <= 250 W)					HPR SLVR fan (CPU TDP up to 350 W)	STD fan (CPU TDP <= 165 W)	HPR GOLD fan (CPU TDP up to 350 W)	HPR SLVR fan 70% (CPU TDP up to 165 W)^	
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

NOTE: In 12 x 3.5-inch with rear module configuration, for CPU TDP greater than 270 W and specific Low Temperature-case CPUs are not supported.

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

Table 46. Thermal restriction matrix for rear NVMe drives with air cooled configuration (non-GPU)

Configuration			24 x 2.5-inch SAS		12 x 3.5-inch	
Rear storage			2 x 2.5-inch with rear fan	4 x 2.5-inch with rear fan	2 x 2.5-inch with rear fan	4 x 2.5-inch with rear fan
Drive type	Drives capacity	Power	HPR SLVR fan		HPR GOLD fan 70%	
Kioxia CD7	15.36 TB	19 W	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
Samsung PM9A3	7.68 TB	14 W	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
Samsung PM1733	15.36 TB	22 W	30°C (86°F)	30°C (86°F)	N/A	N/A
Samsung PM1733a	15.36 TB	19.7 W	35°C (95°F)	30°C (86°F)	30°C (86°F)	N/A
Samsung PM1735a	12.8 TB	19.8 W	35°C (95°F)	30°C (86°F)	30°C (86°F)	N/A
Redtail	7.68 TB	24.5 W	30°C (86°F)	30°C (86°F)	N/A	N/A
Hynix PE8010	7.68/3.84/1.92 TB	17 W	35°C (95°F)	30°C (86°F)	30°C (86°F)	N/A
Intel P5520	15.36 TB	20 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
Kioxia CM7	30.72 TB	25 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
Kioxia CD8	15.36 TB	19 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
PE8110	7.68 TB	20 W	30°C (86°F)	N/A	N/A	N/A
PE8110	3.84/1.92 TB	20 W	35°C (95°F)	30°C (86°F)	30°C (86°F)	N/A
PS1010	15.36 TB	20 W	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)

Table 46. Thermal restriction matrix for rear NVMe drives with air cooled configuration (non-GPU) (continued)

Configuration			24 x 2.5-inch SAS		12 x 3.5-inch	
Rear storage			2 x 2.5-inch with rear fan	4 x 2.5-inch with rear fan	2 x 2.5-inch with rear fan	4 x 2.5-inch with rear fan
Drive type	Drives capacity	Power	HPR SLVR fan		HPR GOLD fan 70%	
PS1030	12.8 TB	20 W	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)

Table 47. Thermal restriction matrix for GPU configurations

Configuration				No back plane	8 x 2.5-inch NVMe	8 x 2.5-inch NVMe + 8 x 2.5-inch SAS	16 x 2.5-inch SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives
CPU TDP/cTDP		Cores	T-Case max center (°C)	HPR GOLD fan with 1U HPR L-Type HSK							
3408U	125 W ¹	8	79	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
5415+	150 W ¹	8	78	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
4410Y		12	78								
5416S		16	78								
5418N	165 W ¹	24	84	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
5411N		24	84								
4416+		20	82								
6426Y	185 W ¹	16	72	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
5418Y		24	80								
5412U		24	80								
6428N		32	85								
6421N		32	85								
6434	205 W ¹	8	96	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (86°F)	30°C (86°F)
5420+		28	84								
6438Y+		32	76								
6438M		32	84								
6438N		32	84								
6442Y	225 W ¹	24	79	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
6448Y		32	79								
6444Y	270 W ²	32	75	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

Table 47. Thermal restriction matrix for GPU configurations (continued)

Configuration				No back plane	8 x 2.5-inch NVMe	8 x 2.5-inch NVMe + 8 x 2.5-inch SAS	16 x 2.5-inch SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives
CPU TDP/cTDP		Cores	T-Case max center (°C)	HPR GOLD fan with 1U HPR L-Type HSK							
8462Y+	300 W ²	32	81	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)
6458Q	350 W ²	32	64	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC
6414U	250 W ²	32	76	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
6454S	270 W ²	32	71	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)
6430		32	71								
8471N	300 W ²	52	76	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)
8470N		52	76								
8460Y+		40	75								
8452Y		36	75								
8480+	350 W ²	56	79	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	Required DLC	Required DLC	Required DLC
8470		52	79								
8468		48	79								
8470Q	350 W ²	52	57	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC
9480	350 W ²	56	64	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC	Required DLC
9470		52	64								
9460	350 W ²	40	77	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	Required DLC	Required DLC	Required DLC
9462		32	77								

NOTE: The platform supports Maximum (MAX) and Mainstream (MS) system boards.

- ¹ supports MS system board (CPU TDP < 250 W)
- ² supports MAX system board (CPU TDP => 250 W)

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

NOTE: *Supported ambient temperature is 30°C (86°F).

NOTE: GPU configuration supports only High performance Gold (HPR Gold) fan.

Table 48. Thermal restriction matrix for memory with air cooled configuration (GPU)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS *	16 x 2.5-inch or 16 x EDSFF E3.S NVMe **	24 x 2.5-inch SAS*	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe***	24 x 2.5-inch NVMe***
DIMM Configuration	2DPC/Power	HPR GOLD fan with 1U HPR L-Type HSK						
256 GB RDIMM	12.7 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	Required DLC	Required DLC	Required DLC
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

NOTE: *In 16 x 2.5-inch SAS and 8 x 2.5-inch NVMe configurations, for CPU TDP 350 W supported ambient temperature is 30°C (86°F).

NOTE: **In 16 x 2.5-inch NVMe configuration, for CPU TDP greater than 300 W supported ambient temperature is 30°C (86°F).

NOTE: ***In 24 x 2.5-inch SAS/NVMe configuration and 16 x 2.5-inch SAS + 8 x 2.5-inch NVMe, for CPU TDP 270 W - 300 W and specific Low Temperature-case CPUs supported ambient temperature is 30°C (86°F).

Table 49. Optimized Ecological upgrade thermal restriction matrix for air cooled configuration

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch			Ambient temperature	
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch rear drives with rear fan	EDSFF E3.S rear drives with rear fan		
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan/HSK								HPR GOLD fan 70%^				
5415+	150 W	8	78	STD /2U HPR	STD /2U HPR	STD /2U HPR	STD /2U HPR	STD /2U HPR	HPR SLVR /2U HPR	STD /2U HPR	HPR GOL	HPR SLVR	HPR GOL D /	HPR SLVR /2U HPR		
4410Y		12	78													

Table 49. Optimized Ecological upgrade thermal restriction matrix for air cooled configuration (continued)

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch			Ambient temperature	
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch rear drives with rear fan	EDSFF E3.S rear drives with rear fan		
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan/HSK								HPR GOLD fan 70%^				
5416S		16	78								D / STD	/2U HPR	2U HPR			
5418N/5411N	165 W	24	84	STD /2U HPR	STD /2U HPR	STD /2U HPR	STD /2U HPR	STD /2U HPR	HPR SLVR /2U HPR	STD /2U HPR	HPR GOLD / STD	HPR SLVR /2U HPR	HPR GOLD /2U HPR	HPR SLVR /2U HPR		
4416+		20	82													

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

Thermal restriction matrix for 5th Gen Intel® Xeon® Scalable Processors

Table 50. Thermal restriction matrix for air cooled configuration

Configuration				No back plane	8 x 2.5-inch NV Me	16 x 2.5-inch SAS and Split NVM e-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVM e	24 x 2.5-inch NV Me	12 x 3.5-inch^		Ambient temperature
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
CPU TDP/cTDP		Cores	T-Case max center (°C)	Fan										
4509Y	125 W ¹	8	84	STD	STD	STD	STD	STD	HPR SLVR	STD	HPR GOLD	HPR SLVR	HPR GOLD	35°C (95°F)
4510	150 W ¹	12	84	STD	STD	STD	STD	STD	HPR SLVR	STD	HPR GOLD	HPR SLVR	HPR GOLD	35°C (95°F)
4514Y		16	79											
5512U	185 W ¹	28	89	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD	HPR GOLD	35°C (95°F)
6534	195 W ¹	8	64	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD	HPR GOLD	35°C (95°F)
6526Y		16	82											
6542Y	250 W ¹	24	83	STD	STD	STD	STD	STD	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD*	HPR GOLD*	35°C (95°F)
6548Y+		32	83											
6548N		32	83											
8562Y+	300 W ²	32	81	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD	Required DLC	Required DLC	35°C (95°F)
8558U	300 W ²	48	78	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD	Required DLC	Required DLC	35°C (95°F)
8568Y+	350 W ²	48	81	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR fan	HPR SLVR fan	HPR GOLD*	Required DLC	Required DLC	35°C (95°F)
8580		60	81											
8592+		64	81											

NOTE: The platform supports Maximum (MAX) and Mainstream (MS) system boards.

- ¹ supports MS system board (CPU TDP < 250 W)
- ² supports MAX system board (CPU TDP ≥ 250 W)

i **NOTE:** *Supported ambient temperature is 30°C (86°F).

i **NOTE:** ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

Table 51. Thermal restriction matrix for memory with air cooled configuration (non-GPU)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch	
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan
DIMM Configuration	2DP C/Power	STD fan (CPU TDP <= 250 W)						HPR SLVR fan (CPU TDP up to 350 W)	STD fan (CPU TDP <= 165 W)	HPR GOLD fan (CPU TDP up to 350 W)	HPR SLVR fan 70% (CPU TDP up to 165 W)^
256 GB RDIMM*	12.7 W	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	35°C (95°F)	Required DLC	35°C (95°F)	Required DLC	Required DLC
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
96 GB RDIMM	8.3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
DIMM Configuration	2DP C/Power	HPR SLVR fan (CPU TDP up to 350 W)							HPR GOLD fan (CPU TDP up to 350 W)	HPR GOLD fan 70% (CPU TDP up to 250 W)^	
256 GB RDIMM*	12.7 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	Required DLC	Required DLC
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)

Table 51. Thermal restriction matrix for memory with air cooled configuration (non-GPU) (continued)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch	
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan
DIMM Configuration	2DP C/Power	STD fan (CPU TDP <= 250 W)						HPR SLVR fan (CPU TDP up to 350 W)	STD fan (CPU TDP <= 165 W)	HPR GOLD fan (CPU TDP up to 350 W)	HPR SLVR fan 70% (CPU TDP up to 165 W)^
96 GB RDIMM	8.3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

NOTE: *256 GB RDIMM with 5th Gen Intel® Xeon® Scalable Processors will be supported in the future release.

Table 52. Supported ambient temperature for processors with GPU

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives
CPU TDP/cTDP		Cores	T-Case max center (°C)	Support HPR GOLD fan with 1U HPR L-Type HSK						
4509Y	125 W ¹	8	84	35°C	35°C	35°C	35°C	35°C	35°C	35°C
4510	150 W ¹	12	84	35°C	35°C	35°C	35°C	35°C	35°C	35°C
4514Y		16	79							
5512U	185 W ¹	28	89	35°C	35°C	35°C	35°C	35°C	35°C	35°C

Table 52. Supported ambient temperature for processors with GPU (continued)

Configuration				No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
Rear storage				No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives
CPU TDP/cTDP		Cores	T-Case max center (°C)	Support HPR GOLD fan with 1U HPR L-Type HSK						
6534	195 W ¹	8	64	35°C	35°C	35°C	35°C	35°C	35°C	35°C
6526Y		16	82							
6542Y	250 W ¹	24	83	35°C	35°C	35°C	35°C	35°C	35°C	35°C
6548Y+		32	83							
6548N		32	83							
8562Y+	300 W ²	32	81	35°C	35°C	35°C	35°C	30°C	30°C	30°C
8558U	300 W ²	48	78	35°C	35°C	35°C	35°C	30°C	30°C	30°C
8568Y+	350 W ²	48	81	30°C	30°C	30°C	30°C	Required DLC	Required DLC	Required DLC
8580		60	81							
8592+		64	81							

NOTE: The platform supports Maximum (MAX) and Mainstream (MS) system boards.

- ¹ supports MS system board (CPU TDP < 250 W)
- ² supports MAX system board (CPU TDP ≥ 250 W)

NOTE: *Supported ambient temperature is 30°C (86°F).

Table 53. Thermal restriction matrix for memory with air cooled configuration (GPU)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
DIMM Configuration	2DPC/Power	HPR GOLD fan with 1U HPR L-Type HSK						
256 GB RDIMM*	12.7 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	Required DLC	Required DLC	Required DLC
128 GB RDIMM	8.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
96 GB RDIMM	8.3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
64 GB RDIMM	6.9 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
32 GB RDIMM	4.1 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
16 GB RDIMM	3 W	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

 **NOTE:** *256 GB RDIMM with 5th Gen Intel® Xeon® Scalable Processors will be supported in the future release.

Common thermal restrictions for both 4th and 5th Gen Intel processors

Table 54. GPU types support thermal restriction for both air cooling and liquid cooling configuration

Configuration	No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and split NVMe-SAS	16 x 2.5-inch NVMe or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe
Rear storage	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives
GPU	HPR GOLD fan with 1U HPR L-Type HSK						
A40 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)
Intel PVC (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	30°C (86°F)	30°C (86°F)	30°C (86°F)
A100 80 GB (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
A16 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
A30 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
A2 (Max 6)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
H100 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
A800 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
L4 (Max 6)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
L40 (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
Intel ASM (Max 6)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)
L40S (Max 2)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)	35°C (95°F)

Table 55. Thermal restriction for memory with liquid cooled configuration(non-GPU)

Configuration		No backplane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe -SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS		16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	12 x 3.5-inch^		Ambient temperature
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	No rear drives	No rear drives	No rear drives	2.5-inch or EDSFF E3.S rear drives with rear fan	
DIMM Configuration	Power	Fan										
256 GB RDIMM*	12.7 W	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR SLVR	HPR GOLD	HPR GOLD fan	
128 GB RDIMM	8.9 W	STD	STD	STD	STD	STD	STD	STD	STD	HPR GOLD	HPR GOLD	35°C (95°F)
96 GB RDIMM	8.3 W	STD	STD	STD	STD	STD	STD	STD	STD	HPR GOLD	HPR GOLD	35°C (95°F)
64 GB RDIMM	6.9 W	STD	STD	STD	STD	STD	STD	STD	STD	HPR GOLD	HPR GOLD	35°C (95°F)
32 GB RDIMM	4.1 W	STD	STD	STD	STD	STD	STD	STD	STD	HPR GOLD	HPR GOLD	35°C (95°F)
16 GB RDIMM	3 W	STD	STD	STD	STD	STD	STD	STD	STD	HPR GOLD	HPR GOLD	35°C (95°F)

NOTE: ^The fan speed in the 3.5-inch chassis is limited to 70% due to the drive dynamic profile.

NOTE: *256 GB RDIMM with 5th Gen Intel® Xeon® Scalable Processors will be supported in the future release.

Table 56. Thermal restriction for memory with liquid cooled configuration(GPU)

Configuration		No back lane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	Ambient temperature
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	
DIMM Configuration	Power	Fan							
256 GB RDIMM *	12.7 W	HPR GOLD fan							35°C (95°F)
128 GB RDIMM	8.9 W								
96 GB RDIMM	8.3 W								

Table 56. Thermal restriction for memory with liquid cooled configuration(GPU) (continued)

Configuration		No back lane	8 x 2.5-inch NVMe	16 x 2.5-inch SAS and Split NVMe-SAS	16 x 2.5-inch or 16 x EDSFF E3.S NVMe	24 x 2.5-inch SAS	16 x 2.5-inch SAS + 8 x 2.5-inch NVMe	24 x 2.5-inch NVMe	Ambient temperature
Rear storage		No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	No rear drives	
DIMM Configuration	Power	Fan							
64 GB RDIMM	6.9 W								
32 GB RDIMM	4.1 W								
16 GB RDIMM	3 W								

 **NOTE:** *256 GB RDIMM with 5th Gen Intel® Xeon® Scalable Processors will be supported in the future release.

Thermal air restrictions

Table 57. Air cooling configurations thermal restriction for AHSRAE A3 and A4

ASHRAE	A3/40°C (104°F)	A4/45°C (113°F)
PSU	Two PSUs are required in redundant mode. If there is PSU failure, system performance may be reduced.	
PCIe card	Non-Dell qualified peripheral cards and peripheral cards greater than 25 W are not supported.	
GPU/FPGA	Not supported	
DIMM	128 GB, or greater capacity DIMMs are not supported.	
PCIe SSD	Not supported	
Front storage	Not supported in 12 x 3.5-inch SAS configuration.	
Rear storage	Not supported	
Fan	HPR SLVR fans are required.	
Processor	≤ 165 W	≤ 125 W
OCP	Supported with 85°C (185°F) active optic cable.	Supported with 85°C (185°F) active optic cable and cards tier ≤4.
BOSS	BOSS-N1 is supported.	BOS-N1 is not supported.

Table 58. Liquid cooling configurations thermal restriction for AHSRAE A3 and A4

ASHRAE	A3/40°C (104°F)	A4/45°C (113°F)
PSU	Two PSUs are required in redundant mode. If there is PSU failure, system performance may be reduced.	
PCIe card	Non-Dell qualified peripheral cards and peripheral cards greater than 25 W are not supported.	
GPU/FPGA	Not supported	
DIMM	128 GB, or greater capacity DIMMs are not supported.	
PCIe SSD	Not supported	
Front storage	Not supported in 12 x 3.5-inch SAS configuration.	

Table 58. Liquid cooling configurations thermal restriction for AHSRAE A3 and A4 (continued)

ASHRAE	A3/40°C (104°F)	A4/45°C (113°F)
Rear storage	Not supported	
Fan	HPR SLVR fans are required in 2.5-inch configurations systems.	
OCP	Supported with 85°C (185°F) active optic cable.	Supported with 85°C (185°F) active optic cable and cards tier ≤4.
BOSS	BOSS-N1 is supported.	BOSS-N1 is not supported.

Appendix A. Standards compliance

The system conforms to the following industry standards.

Table 59. Industry standard documents

Standard	URL for information and specifications
ACPI Advance Configuration and Power Interface Specification, v6.4	Uefi specifications and tools
Ethernet IEEE Std 802.3-2022	ieee standards
MSFT WHQL Microsoft Windows Hardware Quality Labs	microsoft.com/whdc/system/platform/pcdesign/desguide/serverdg.msp
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi
DDR5 Memory DDR5 SDRAM Specification	jedec.org/standards-documents/docs/jesd79-4.pdf
PCI Express PCI Express Base Specification, v5.0	pcisig.com/specifications/pciexpress
PMBus Power System Management Protocol Specification, v1.2	pmbus specification and revisions
SAS Serial Attached SCSI, 3 (SAS-3) (T10/INCITS 519)	SCSI storage interfaces information
SATA Serial ATA Rev. 3.3	sata-io.org page
SMBIOS System Management BIOS Reference Specification, v3.3.0	BIOS reference specification page
TPM Trusted Platform Module Specification, v1.2 and v2.0	trustedcomputinggroup.org page
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFIF specifications
PI Platform Initialization Specification, v1.7	
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB Implementers Forum, Inc. 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 NVMe Express Management Interface. Revision 1.2c	
NVMe NVMe Boot Specification. Revision 1.0	

Appendix C Additional resources

Table 60. 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 system 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
QR code for system resources	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

Dell PowerEdge R760 NEBS level 3

As IT and Operational Technologies converge, communications service providers are evaluating their infrastructures to meet the demands of digital transformation. Dell is uniquely positioned to assist both service providers and network equipment providers (NEPs) in succeeding. By working closely with customers to fully understand their needs, Dell offers global partnership and collaboration.

- Network Equipment-Building System (NEBS) Level 3 and ETSI validated
- Commercial off the shelf hardware
- Comprehensive, global availability, service, and support

Telecom server solution

NEBS compliance is an important requirement in your environment. Dell Solutions, a global leader in enterprise platforms, provides NEBS Level-3 (GR-63 and GR-1089) and ETSI compliant PowerEdge Servers with Intel® Xeon® Processors, ensuring top stability and global availability. Rack-mount systems are designed to deliver high performance, maximum scalability, and safe and reliable service.

Dell PowerEdge Servers offer:

1. Open standard systems:
 - Improved compatibility based on industry-leading Dell products
 - Rapidly scalable and expandable
2. Industry standards Solutions
 - Validated to NEBS Level 3 standards, VZ.TPR.9205, and various ETSI standards
 - Enabled for operation in warmer environments than traditional data centers – Designed for extreme conditions such as high humidity, earthquakes, and dust
3. Global regulatory support and availability

Following are the additional NEBS validated technology devices available:

- Intel i350 Quad Port 1 GbE BASE-T
- OCP NIC 3.0 [540-BCOE]
- Intel Ethernet i350 Quad Port 1 GbE BASE-T Adapter, PCIe Full Height, V2, FIRMWARE RESTRICTIONS APPLY [540-BDIW]
- Intel E810-XXVDA4 Quad Port 10/25GbE SFP28 Adapter, PCIe Full Height [540-BDDR]

NOTE: We cannot include the Serial Com card as it has not undergone testing on any platforms. Therefore, we cannot confirm it as an additional card.

Service and support

Bring game-changing innovations to market quickly with services including design, manufacturing, fulfillment, and global support. Refine products or design new ones with the right services, allowing focus on IP. Choose from services that can help:

- Conduct applications testing on the hardware
- Integrate hardware, images, applications, peripheral, and documents as your systems are built
- Consolidate, stage, deliver, and support your orders globally
- Deliver Customer Support anywhere with over 30,000 Dell employees in tech support, parts, and field services across 100+ countries

More Features

- Combine density, performance, and scalability to optimize application performance
- Manage your clients more efficiently with industry-leading support

- Ensure server security from the factory to you
 - Rely on a secure component supply chain to ensure protection from the factory to the data center
 - Maintain data safety with cryptographically signed firmware packages and Secure Boot
 - Prevent unauthorized or malicious changes with Server Lockdown
 - Wipe all data from nonvolatile media including hard drives, SSDs, and system memory quickly and securely with SystemErase

Table 61. Specifications

Feature	Technical Specification	
Processor: Available for the entire life cycle. Up to 2x 185W Intel Xeon SP	• Intel Xeon 6428N • Intel Xeon 5412U • Intel Xeon 5418N • Intel Xeon 5411N • Intel Xeon 4514Y • Intel Xeon 4410Y • Intel Xeon 6421N • Intel Xeon 4509Y • Intel Xeon 5415+ • Intel Xeon 5512U • Intel Xeon 3408U • Intel Xeon 6426Y • Intel Xeon 5416S • Intel Xeon 4416+ • Intel Xeon 5418Y • Intel Xeon 4510	
Memory	32 DDR4 DIMM slots supporting RDIMM, speeds up to 4800 8GB, 16GB, 32GB, 64GB capacities supported	
Storage controllers	Front Controllers: PERC H965i, PERC H755	
Drive bays	R760 Front Drive Bays: Up to 24 x 2.5 inch SAS or SATA SSD including 8 Universal Slots (NVMe direct)	
Power supplies	Titanium redundant hot swap 1100 W DC, 1400 W AC, and 1800 W AC power supply supported.	
Sizing	Form factor: 1. R760: Rack (2U)	Chassis depth: 1. R760: 772.13 mm with Bezel
Embedded management	iDRAC9 with Lifecycle Controller (Express, Enterprise)	
Bezel	NEBS Filtered, Dell branded	
OpenManage Software	• OpenManage Enterprise • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin • CloudIQ for PowerEdge plug in • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center • OpenManage Integration with Windows Admin Center	
Integrations and connections	• BMC Truesight • Microsoft System Center • OpenManage Integration with ServiceNow	• IBM Tivoli Netcool/OMNIBus • IBM Tivoli Network Manager IP Edition • Micro Focus Operations Manager

Table 61. Specifications (continued)

Feature	Technical Specification	
	- Red Hat Ansible Modules - Terraform Providers - VMware vCenter and vRealize Operations Manager	- Nagios Core - Nagios XI
Security	- Cryptographically signed firmware - Data at Rest Encryption (SEDs with local or external key mgmt) - Secure Boot - Secured Component Verification (Hardware integrity check) - 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	
Embedded NIC	Broadcom 5720 Dual Port 1 GbE LOM	
I/O Adapter Options	Additional card options: - Intel Ethernet 100G 2P E810-C Adapter (FH and LP) - Broadcom NetXtreme-E P2100D BCM57508 2x100G QSFP PCIE Ethernet (LP) - Mellanox ConnectX-6 Dx Dual Port 100 GbE QSFP56 PCIE Adapter (FH and LP) - Mellanox Bluefield 2 DP 25 GbE SFP Crypto DPU (FH) - Broadcom NetXtreme-E P425D BCM57504 4x25G SFP28 PCIE Ethernet (FH) - Intel(R) Ethernet 25G 2P E810-XXV Adapter (LP) - NVIDIA ConnectX-6LX Dual Port 25G GbE SFP28 Network Adapter (LP) - Broadcom BCM57414 25G SFP Dual Port PCIE adapter (LP) - Intel(R) Ethernet 25G 4P E810-XXV OCP - Broadcom BCM57414 25G SFP Dual Port OCP3 Mezz - Intel(R) Ethernet 25G 4P E810-XXV OCP	
Ports	R760: - Front ports: 1x VGA, 1x USB 2.0, 1x dedicated iDRAC Direct Micro-USB. - Rear ports: 1x VGA, 1 x USB 2.0, 1x USB 3.0, 1x Dedicated iDRAC Ethernet.	
Fans	HPR Gold Fan	
Rack rail support	Combo Rails	
Environmental specs (NEBS Level-3 and ETSI)	- Temperature: Continuous operating temperature of -5C to 40C; 96 hour operating excursions from -5C to 55C - Humidity water/kg of dry air: Operating Humidity of 5% to 85% with excursions of 5% to 90%, but not to exceed 0.24 kg : - Altitude: Up to 4000m; -60m to 1800 m; -61m to 1829m at 40C; 1829m–3960m at 30C - Dust: Dust filter rated 80% per ASHRAE Std 52.1 - Seismic: Operational resiliency up to Richter 7.5 seismic event (Zone 4 seismic event) - EMI: Immunity up to 8kV contact or 15kV air discharge - Fire resistance: Constructed from fire-retardant materials designed to contain and extinguish any fires that may occur inside the box.	