

Dell PowerEdge R6625

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
Chassis views.....	10
Front view of the system.....	10
Rear view of the system.....	11
Inside the system.....	11
Chapter 4: Processor.....	14
Processor features.....	14
Supported processors.....	14
Chapter 5: Memory subsystem.....	16
Supported memory.....	16
Chapter 6: Storage.....	18
Storage controllers.....	18
Supported Drives.....	18
HDD feature matrix.....	19
Supported Drives.....	20
Internal storage configuration.....	20
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.....	31
Power.....	31
Power.....	32
Thermal.....	33
Acoustics.....	34
Chapter 10: Rack, rails, and cable management.....	39

Rails and cable management information.....	39
Chapter 11: Supported Operating Systems.....	48
Chapter 12: Dell Systems Management.....	49
Integrated Dell Remote Access Controller (iDRAC).....	49
Systems Management software support matrix.....	50
Chapter 13: Appendix A. Additional specifications.....	52
Chassis dimensions	52
System weight.....	53
NIC port specifications.....	53
Video specifications.....	53
USB ports specifications.....	54
PSU rating.....	54
Environmental specifications.....	55
Particulate and gaseous contamination specifications.....	57
Thermal restriction matrix.....	58
Chapter 14: Appendix B. Standards compliance.....	62
Chapter 15: Appendix C Additional resources.....	63

System overview

The Dell™ PowerEdge™ R6625 (2S1U) is Dell's latest 2-socket, 1U rack servers designed to run complex workloads using highly scalable memory, I/O, and network options.

The system features:

- Two AMD EPYC 4th Generation 9004 series processor with up to 128 cores
- Optional Direct Liquid Cooling (DLC) for required CPUs and/or configurations
- Up to 24 RDIMMs with up to 6 TB of memory and speeds up to 4800 MT/s
- Two redundant AC or DC power supply units
- Up to 10 x 2.5-inch SATA/SAS/NVMe or 8 x 2.5-inch SATA/SAS/NVMe drives or 4 x 3.5-inch SATA/SAS
- Up to 14 x E3.S NVMe direct drives or 16 x E3.S, 2 x E3.S (rear) NVMe direct drives
- PCI Express® (PCIe) 5.0 enabled expansion slots
- Network interface technologies to cover Network Interface Card (NIC)

The PowerEdge R6625 is a general-purpose platform capable of handling demanding workloads and applications, such as data warehouses, eCommerce, databases, and high-performance computing (HPC).

Topics:

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

Key workloads

Customers looking for accelerated compute to maximize performance in a dense, scalable server architecture to address the following applications:

- High Performance Computing
- Virtual Desktop Infrastructure (VDI)
- Virtualization

New technologies

The PowerEdge R6625 is capable of handling demanding workloads and applications, such as data warehouses, e-commerce, databases, and high-performance computing (HPC).

Table 1. New technologies

Technology	Detailed Description
AMD Genoa Processor (SP5)	Core count: Up to 128 core processor
	5nm process technology
	AMD Inter-chip global memory interconnect (xGMI) up to 64 lanes
	Speeds up to 4.1 GHz
	Maximum TDP: 400 W
5600 MT/s DDR5 Memory  NOTE: The speed of 5600 MT/s RDIMM will be clocked down to match the CPU MAX DDR frequency (1 DPC) at 4800 MT/s supported on the platform.	Up to 12 channels with 1 DPC per CPU and 24 DIMMs in total
	Supports RDIMM, 3DS DIMM DDR5 with ECC up to 4800 MT/s

Table 1. New technologies (continued)

Technology	Detailed Description
PCIe Gen	Gen5 @32 GT/s
PCIe Slot	Up to three PCIe Slots with x16 lanes in total
Flex I/O	LOM board, 2 x1Gb with BCM5720 LAN controller
	Rear I/O with: • 1 x Dedicated iDRAC ethernet port • 1 x USB 3.0 • 1 x USB 2.0 • 1 x VGA
	Serial Port Option with STD RIO board
	OCP Mezz 3.0 (supported by x8 PCIe lanes) (optional)
	Front I/O with: • 1 x USB 2.0 • 1 x VGA • 1 x iDRAC Direct (Micro-AB USB) port
CPLD 1-wire	Support payload data of Front PERC, Riser, BOSS N1, BP, PUCK, and Rear IO to BIOS and iDRAC
Dedicated PERC	PERC 11 • HBA355i, H355, H755, H755N
	PERC 12 • H965i • H965e • HBA465i • HBA465e
Software RAID	OS RAID / S160 NVME SW RAID
Power Supplies	60 mm dimension is the new PSU form factor on 16G design
	Titanium 1800 W AC/HVDC
	Platinum 1400 W AC/HVDC
	Titanium 1400 W AC/HVDC
	Titanium 1100 W AC/HVDC
	Platinum 800 W AC/HVDC

System features and generational comparison

The following table shows the comparison between the PowerEdge R6625 with the PowerEdge R6525.

Table 2. Features comparison

Features	PowerEdge R6625 (2S1U)	PowerEdge R6525 (2S1U)
Processors	Two AMD® EPYC 4 th Generation Genoa (SP5) processors	Two AMD® EPYC™ Generation 2 or Generation 3 processors.
Processor interconnect	Inter-chip global memory interconnect (xGMI) 32GT/sec.	Inter-chip global memory interconnect (xGMI) 16GT/sec.
Memory	24x DDR5 RDIMM, 3DS DIMM DDR5 with ECC, bandwidth up to 5600 MT/s i NOTE: The speed of 5600 MT/s RDIMM will be clocked down to match the CPU MAX DDR frequency (1 DPC) at 4800 MT/s supported on the platform.	32x DDR4 RDIMM, LRDIMM, 3DS
Storage Controllers	- PERC 11: HBA355i, H355, H755, H755N - PERC 12: H965i, H965e, HBA465e, HBA465i - Software RAID: S160 - BOSS-N1	- PERC: H755N, H840, H745, HBA345, H345, HBA355, HBA355e - Software RAID: S150 - BOSS-S1 - BOSS-S2
PCIe SSD	- Up to 12 (10+2) NVMe direct attach. - Up to 16 x E.3 Gen.5	Up to 10+2 PCIe SSD
Drive Bays	Front bays: - Up to 4 x 3.5-inch SAS/SATA (HDD/SSD) max 80 TB - Up to 8 x 2.5-inch NVMe max 122.88 TB - Up to 10 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 153.6 TB - Up to 14 x E3.S (NVMe Gen5) max 89.6 TB - Up to 16 x E3.S (NVMe Gen5) max 102.4 TB Rear bays: - Up to 2 x 2.5-inch SAS/SATA (HDD/SSD) max 30.72 TB - Up to 2 x E3.s (NVMe Gen5) max 15.36 TB	Front bays: - Up to 4 x 3.5-inch SAS/SATA drives - Up to 8 x 2.5-inch SAS/SATA drives - Up to 10 x 2.5-inch SAS/SATA/NVMe drives Rear bays: - Up to 2 x 2.5-inch SAS/SATA/NVMe drives
Power Supplies	1800 W Mixed Mode Titanium 100-240 VAC or 240 HVDC 1400 W Mixed Mode Titanium 100-240 VAC or 240 HVDC 1400 W Mixed Mode Platinum 100-240 VAC or 240 HVDC 1400 W Mixed Mode Titanium 277 VAC or 336 HVDC 1100 W Mixed Mode Titanium 100-240 VAC or 240 HVDC	1400 W Mixed Mode 100-240 VAC or 240 HVDC 1100 W Mixed Mode 100-240 VAC or 240 HVDC 1100 W LVDC -48 - -60 VDC 800 W Platinum 100-240 VAC or 240 HVDC Hot swap PSUs with full redundancy.

Table 2. Features comparison (continued)

Features	PowerEdge R6625 (2S1U)	PowerEdge R6525 (2S1U)	
	1100 W LVDC -48 - -60 VDC, hot swap redundant 800 W Platinum 100-240 VAC or 240 HVDC Hot swap PSUs with full redundancy.		
Cooling Options	- Air Cooling - Optional Direct Liquid Cooling (DLC)	Air Cooling	
Fans	Up to four sets (dual fan module) Standard (STD) / High-Performance Gold (HPR Gold) hot plug fans	Up to four Standard (STD)/ High-performance silver grade (HPR Gold (Silver)), or high performance gold grade (HPR Gold (Gold))	
Dimension	Height - 42.8 mm (1.685 inches)	Height - 42.8 mm (1.685 inches)	
	Width - 482.0 mm (18.97 inches)	Width - 482.0 mm (18.97 inches)	
	Depth - 822.89 mm (32.4 inches) with bezel	Depth - 822.89 mm (32.4 inches) with bezel	
	Depth - 809.05 (31.85 inches) without bezel	Depth - 809.05 (31.85 inches) without bezel	
Form Factor	1U rack server	1U rack server	
Embedded Management	- iDRAC9 - iDRAC Direct - iDRAC RESTful API with Redfish - iDRAC Service Manual - Quick Sync 2 wireless module	- iDRAC9 - iDRAC Direct - iDRAC RESTful API with Redfish - iDRAC Service Manual - 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 plug-in - OpenManage Services plug-in - OpenManage Update Manager plug-in	- OpenManage Enterprise - LC 3.x - OpenManage Power Manager plug-in - OpenManage Services plug-in - OpenManage Update Manager plug-in	
Mobility	OpenManage Mobile	OpenManage Mobile	
Integrations and Connections	OpenManage Integrations - BMC True sight - Microsoft System Center - OpenManage Integration with ServiceNow - Red Hat Ansible Modules - Terraform Providers - VMware vCenter and vRealize Operations Manager	OpenManage Integrations - Microsoft System Center - Windows Admin Center (WAC) - ServiceNow - Red Hat Ansible Modules - VMware vCenter	- OM Integration with Microsoft System Center (OMIMSC) - OM Integration with Vcenter/vROps
Security	- AMD Secure Encrypted Virtualization (SEV) - AMD Secure Memory Encryption (SME) - 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	- Cryptographically signed firmware - Secure Boot - Secure Erase - Silicon Root of Trust - System Lockdown (requires iDRAC9 or Enterprise Datacenter) - TPM 1.2/2.0 (optional), TCM 2.0 optional - AMD Secure Memory Encryption (SME) - AMD Secure Encrypted Virtualization (SEV)	

Table 2. Features comparison (continued)

Features	PowerEdge R6625 (2S1U)		PowerEdge R6525 (2S1U)	
	- System Lockdown (requires iDRAC9 Enterprise or Datacenter) - TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ			
Embedded NIC	2 x 1 GbE LOM (optional)		No LOM riser supported on R6525.	
Networking Options	1 x OCP card 3.0 (optional) <i>i</i> NOTE: The system allows either LOM card or an OCP card or both to be installed in the system.		OCP 3.0 SCFF (Small Card Form Factor)	
GPU Options	2 x 75 W (SW)		2 x 75 W (SW/FH) 1 x 75 W (SW/FH) 3 x 75 W (SW/LP)	
Ports	Front Ports 1 x iDRAC Direct (Micro-AB USB) port 1 x USB 2.0 1 x VGA	Rear Ports 1 x Dedicated iDRAC Ethernet port 1 x USB 2.0 1 x USB 3.0 1 x Serial (optional) 1 x VGA (optional for Direct Liquid Cooling configuration)	Front Ports 1 x iDRAC USB (Micro-AB USB) 1 x USB 2.0	Rear Ports 1 x USB 3.0 2 x USB 2.0
	Internal Port: 1 x USB 3.0		Internal Port: 1 x USB 3.0	
PCIe	Up to three PCIe slots on Riser cards 2 x PCIe Gen5 slots 3 x PCIe Gen4 slots		Up to three PCIe slots (4.0 x16)	
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 - Red Hat Enterprise Linux - SUSE Linux Enterprise Server - VMware ESXi - Cent operating system For specifications and interoperability details, see Dell Enterprise Operating Systems on Servers, Storage, and Networking page at Dell.com/OSsupport .	

Chassis views and features

Topics:

- Chassis views

Chassis views

Front view of the system



Figure 1. Front view of 4 x 3.5-inch drive system



Figure 2. Front view of 8 x 2.5-inch drive system



Figure 3. Front view of 10 x 2.5-inch drive system



Figure 4. Front view of 14 x E3.S drive system



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

Rear view of the system



Figure 6. Rear view of the system with no risers



Figure 7. Rear view of the system with 2 x 2.5-inch rear drive system



Figure 8. Rear view of the system with 2 x E3.S rear drive system

Inside the system




Figure 9. Inside the system with Rear HDD module and Riser 3




Figure 10. Inside the system with Liquid Cooling




Figure 11. Inside the system with no Risers




Figure 12. Inside the system with Riser 2 and Riser 3 module

Processor

Topics:

- [Processor features](#)

Processor features



The AMD EPYC™9004 Series Processor ("Genoa") is the 4th Generation AMD EPYC™ System on a chip (SOC) supporting the modern data center. The AMD EPYC™ 9004 Series Processor is based on AMD's SP5 compatible socket infrastructure with a new BIOS. AMD EPYC™9004 Series Processor is not drop-in compatible to AMD's SP3 socket infrastructure for EPYC™ Series 7002 ("Rome") and 7003 ("Milan") Processors. Based on AMD's "Zen4" and "Zen4c" cores, integrated I/O controllers, up to 32 MB of L3 cache per core, advanced security, and synchronized fabric and memory clock speeds, "Genoa" is designed for improved performance, lower TCO, and faster time to results with next gen technologies.

The following lists the features and functions included in the AMD Genoa offering:

- Supports up to 128 AMD Zen4c, 96 AMD Zen4 cores with enhance performance with 32 MB L3 cache/core. Integrated I/O support for up to 128 lanes with PCI Express 5.
- Enhanced Memory Performance with support up to 4800 MT/s DIMMs (1 DPC), and RDIMM support. Enhanced Memory Performance with: Infinity Fabric™ and Memory Clock Synchronized. Largest Available x86 L3 Cache –Up to 32 MB/core.
- Memory Capacity with up to 12 channels DDR5 and up to 256 GB/channel support with 2, 4, 6, 8, 10 and 12 channel performance optimization options.
- Enhanced physical and virtual security with AMD Infinity Guard that includes silicon embedded security and virtual features (Secure Memory Encryption and Secure Encrypted Virtualization-Secure Nested Paging (SEV-SNP)to further improve platform and data security capabilities.

Supported processors

Table 3. Supported processor list for the PowerEdge R6625

Processor model number	Default TDP (W)	cTDP Min (W)	cTDP Max (W)	Cores	Threads	L3 Cache (MB)	Base frequency in GHz
9754	360	360	400	128	256	256	2.25
9734	340	340	400	112	224	256	2.20
9684X	320	320	400	96	192	1152	2.55
9654	360	320	400	96	192	384	2.40
9634	290	240	300	84	168	384	2.25

Table 3. Supported processor list for the PowerEdge R6625 (continued)

Processor model number	Default TDP (W)	cTDP Min (W)	cTDP Max (W)	Cores	Threads	L3 Cache (MB)	Base frequency in GHz
9554	360	320	400	64	128	256	3.10
9534	280	240	300	64	128	256	2.45
9454	290	240	300	48	96	256	2.75
9384X	320	320	400	32	64	768	3.10
9354	280	240	300	32	64	256	3.25
9334	210	200	240	32	64	128	2.70
9254	200	200	240	24	48	128	2.9
9224	200	200	240	24	48	384	2.5
9124	200	200	240	16	32	64	3.00
9474F	360	320	400	48	96	256	3.6
9374F	320	320	400	32	64	256	3.85
9274F	320	320	400	24	48	256	4.05
9184X	320	320	400	16	32	768	3.55
9174F	320	320	400	16	32	256	4.10

 **NOTE:** The processors are 12 channel and has a maximum frequency of 4800 MT/s (1DPC).

Memory subsystem

Topics:

- Supported memory

Supported memory

The R6625 supports up to 24 DIMMs (12 per socket), with up to 6 TB of memory and speeds of up to 4800 MT/s.

The R6625 support registered (RDIMMs) which use a buffer to reduce memory loading and provide greater density, allowing for the maximum platform memory capacity. Unbuffered DIMMs (UDIMMs) are not supported.

Table 4. Memory technology comparison

Feature	PowerEdge R6625 (DDR5)
DIMM type	RDIMM
Transfer speed	5600 MT/s
Voltage	1.1 V

NOTE: The speed of 5600 MT/s RDIMM will be clocked down to match the CPU MAX DDR frequency (1 DPC) at 4800 MT/s supported on the platform.

The following table lists the supported DIMMs for the R6625.

Table 5. Supported memory matrix

DIMM type	Rank	Capacity	Operating Speed
			1 DIMM per channel (DPC)
DDR5 RDIMM	1 R	16 GB	4800 MT/s
	2 R	32 GB, 64 GB	4800 MT/s
	2 R	96 GB	4800 MT/s
	2 R	128 GB	5600 MT/s
	4 R	128 GB	4800 MT/s
	8 R	256 GB	4800 MT/s

Table 6. Supported DIMMs

DIMM Speed (MT/s)	DIMM Type	DIMM Capacity (GB)	Ranks per DIMM	Data Width	DIMM Volts (V)
4800	RDIMM	16	1R	x8	1.1
4800	RDIMM	32	2R	x8	1.1
4800	RDIMM	64	2R	x4	1.1
5600	RDIMM	96	2R	x4	1.1
4800	RDIMM	128	4R	x4	1.1
4800	RDIMM	256	8R	x4	1.1

Table 6. Supported DIMMs (continued)

DIMM Speed (MT/s)	DIMM Type	DIMM Capacity (GB)	Ranks per DIMM	Data Width	DIMM Volts (V)
5600	RDIMM	16	1R	x8	1.1
5600	RDIMM	32	2R	x8	1.1
5600	RDIMM	64	2R	x4	1.1
5600	RDIMM	128	2R	x4	1.1

 **NOTE:** 256 GB RDIMM is not supported in 4 x 3.5-inch chassis configuration.

Storage

Topics:

- [Storage controllers](#)
- [Supported Drives](#)
- [Internal storage configuration](#)
- [External Storage](#)

Storage controllers

Dell's 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 Harpoon-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 (SATA, NVMe), SW RAID SATA, NVMe
Value	H355, HBA355 (internal/External), HBA465 (internal/external)
Premium Performance	H755, H755N, H965 (internal/external)

 **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](#).

Supported Drives

Table 8. Supported Drives

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
2.5 inches	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
3.5 inches	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
2.5 inches	vSAS SSD	12 Gb	N/A	960 GB, 1.92 TB, 3.84 TB, 7.68 TB
2.5 inches	SAS SSD	24 Gb	N/A	800 GB, 960 GB, 1.6 TB, 1.92 TB, 3.84 TB, 7.68 TB
2.5 inches	SATA SSD	6 Gb	N/A	480 GB, 960 GB, 1.92 TB, 3.84 TB
2.5 inches	NVMe SSD	Gen4	N/A	400 GB, 800 GB, 1.6 TB, 1.92 TB, 3.2 TB, 3.84 TB, 6.4 TB, 7.68 TB, 15.36 TB

Table 8. Supported Drives (continued)

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	DC NVMe SSD	Gen4	N/A	960 GB, 3.84 TB
EDSFF E3.S	NVMe	Gen5	SSD	3.84 TB, 7.68 TB

HDD feature matrix

Table 9. HDD feature matrix

Type	Interface	Form Factor	RPM	Sector	Security	Capacity
HDD	12Gbps SAS	2.5	10K	512n	ISE	600 GB
HDD	12Gbps SAS	2.5	10K	512n	ISE	1.2 TB
HDD	12Gbps SAS	2.5	10K	512n	FIPS-140	1.2 TB
HDD	12Gbps SAS	2.5	10K	512e	ISE	2.4 TB
HDD	12Gbps SAS	2.5	10K	512e	FIPS-140	2.4 TB
HDD	12Gbps SAS	2.5	7.2K	512n	ISE	2 TB
HDD	12Gbps SAS	2.5	7.2K	512n	ISE	4 TB
HDD	12Gbps SAS	2.5	7.2K	512e	ISE	8 TB
HDD	12Gbps SAS	2.5	7.2K	512e	FIPS-140	8 TB
HDD	12Gbps SAS	2.5	7.2K	512e	ISE	12 TB
HDD	12Gbps SAS	2.5	7.2K	512e	ISE	16 TB
HDD	12Gbps SAS	2.5	7.2K	512e	FIPS-140	16 TB
HDD	12Gbps SAS	2.5	7.2K	512e	ISE	20 TB
HDD	6Gbps SATA	3.5	7.2K	512n	ISE	2 TB
HDD	6Gbps SATA	3.5	7.2K	512n	ISE	4 TB
HDD	6Gbps SATA	3.5	7.2K	512e	ISE	8 TB
HDD	6Gbps SATA	3.5	7.2K	512e	ISE	12 TB
HDD	6Gbps SATA	3.5	7.2K	512e	ISE	16 TB

Table 9. HDD feature matrix (continued)

Type	Interface	Form Factor	RPM	Sector	Security	Capacity
HDD	6Gbps SATA	3.5	7.2K	512e	ISE	20 TB

Supported Drives

Table 10. Supported Drives

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
2.5 inches	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
3.5 inches	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
2.5 inches	vSAS SSD	12 Gb	N/A	960 GB, 1.92 TB, 3.84 TB, 7.68 TB
2.5 inches	SAS SSD	24 Gb	N/A	800 GB, 960 GB, 1.6 TB, 1.92 TB, 3.84 TB, 7.68 TB
2.5 inches	SATA SSD	6 Gb	N/A	480 GB, 960 GB, 1.92 TB, 3.84 TB
2.5 inches	NVMe SSD	Gen4	N/A	400 GB, 800 GB, 1.6 TB, 1.92 TB, 3.2 TB, 3.84 TB, 6.4 TB, 7.68 TB, 15.36 TB
2.5 inches	DC NVMe SSD	Gen4	N/A	960 GB, 3.84 TB
EDSFF E3.S	NVMe	Gen5	SSD	3.84 TB, 7.68 TB

Internal storage configuration

Table 11. Internal Storage Configuration Matrix

Total HDD/SDD (not BOSS)	NVMe Enabled / Universal Slots	15G Storage front	Rear Storage	PERC Quantity	Storage Controller (s)	Controller form factor
0 *BOSS/IDSDM Mandatory	0 / 0	No front backplane	N/A	0	N/A	N/A
4	0 / 0	4 x 3.5-inch SAS/SATA (passive)	N/A	1	SATA S160	N/A
4	0 / 0	4 x 3.5-inch SAS/SATA (passive)	N/A	1	SATA S160	Front PERC
6	0 / 0	4 x 3.5-inch SAS/SATA (passive)	Rear 2 x 2.5-inch Universal (passive)	1	HBA345/HBA355/H345/H745	Front PERC
6	2 / 0	4 x 3.5-inch SAS/SATA (passive)	Rear 2 x 2.5-inch Universal (passive)	1	HBA345/HBA355/H345/H745 S160	Front PERC

Table 11. Internal Storage Configuration Matrix (continued)

Total HDD/ SDD (not BOSS)	NVMe Enabled / Universal Slots	15G Storage front	Rear Storage	PERC Quant ity	Storage Controller (s)	Controller form factor
8	0 / 0	8 x 2.5-inch SAS/SATA (passive)	N/A	1	HBA345/HBA355/H345/ H745	Front PERC
10	0 / 0	10 x 2.5-inch SAS/SATA (Universal)	N/A	1	HBA345/HBA355/H345/ H745	Front PERC
12	0 / 0	10 x 2.5-inch Universal (passive)	N/A	1	HBA345/HBA355/H345/ H745	Front PERC
10	10 / 0	10 x 2.5-inch Universal (passive)	Rear 2 x 2.5- inch Universal (passive)	1	HBA345/HBA355/H345/ H745 S160	Front PERC
10	10 / 0	10 x 2.5-inch Universal (passive)	N/A	1	S160	N/A
12	12 / 0	10 x 2.5-inch Universal (passive)	Rear 2 x 2.5- inch Universal (passive)	1	S160	N/A
8	8 / 0	10 x 2.5-inch Universal (passive)	N/A	1	H755N	Front PERC
14	14 / 0	14 x E3.S G4 Ortho	N/A	2	H755N F	Front PERC
14	14 / 0	14 x E3.S G4 Ortho	N/A	2	H965i F	Front PERC
14	14 / 0	14 x E3.S G5 Ortho	N/A	0	S160	N/A
16	16 / 0	16x E3.S G4 Ortho	N/A	2	H755N F	Front PERC
16	16 / 0	16x E3.S G4 Ortho	N/A	2	H965i F	Front PERC
16	16 / 0	16x E3.S G5 Ortho	N/A	0	S160	N/A

Table 12. PCIe Configurations with PERC and rear storage

Config #	RSR Configuration	# of CPUs	PERC type supported	Rear Storage Possible
0	No RSR	2	Front PERC	No
1	R2A + R3A	2	Front PERC	No
2	R2P + R3P	2	Front PERC	No
3	R1P+R4P	2	Front PERC	No
4	R2Q	2	Front PERC	No
5	R3P	2	Front PERC	Yes
7	R2A	1	Front PERC	No

External Storage

The R6625 support the external storage device types listed in the table below.

Table 13. Support for 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 ME5 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 and Lifecycle Controller. These adapters are rigorously validated for worry-free, fully supported use in our servers.

OCP 3.0 support

Table 14. OCP 3.0 feature list

Feature	OCP 3.0
Form factor	SFF
PCIe Gen	Gen4
Max PCIe width	x8, x16 (with OCP cable)
Max no.of ports	4
Port type	BT/SPF/SFP+/SFP28/SFP56/Q56
Max port speed	25 GbE, 100 GbE (with OCP cable)
NC-SI	Yes
WoL	Yes
Power consumption	15 W–35 W

Supported OCP cards

Table 15. Supported OCP cards

Form Factor	DPN	Vendor	Port type	Max Port speed	Port Count	Max PCIe width
OCP 3.0	2VY37	Broadcom	Q56	100 GbE	2	x16
	FD63G	Mellanox	SFP56	100 GbE	2	x16
	DN78C	Mellanox	SFP28	25 GbE	2	x16
	3Y64D	Broadcom	SFP28	25 GbE	4	x16
	24FG6	Broadcom	SFP28	25 GbE	2	x8
	PWH3C	Intel	SFP28	25 GbE	2	x8
	Y4VV5	Intel	SFP28	25 GbE	4	x16

Table 15. Supported OCP cards (continued)

Form Factor	DPN	Vendor	Port type	Max Port speed	Port Count	Max PCIe width
	W5HC8	Broadcom	BT	10 GbE	4	x16
	F6X1R	Intel	BT	10 GbE	2	x8
	RN1M5	Broadcom	BT	10 GbE	2	x8
	G9XC9	Broadcom	BT	1 GbE	4	x8
	D1C51	Intel	BT	1 GbE	4	x4
	HY4CV	Intel	BT	1 GbE	4	x4
	VJWVJ	Broadcom	BT	1 GbE	4	x4

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 16. OCP 3.0, 2.0, and rNDC NIC comparison

Form Factor	Dell rNDC	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCIe Gen	Gen 3	Gen 3	Gen 4	Supported OCP3 are SFF (small form factor)
Max PCIe Lanes	x8	Up to x16	Up to x16 or x32	See server slot priority matrix
Max of Ports	4	4	4	NA
Port Type	BT/SFP/SFP+/SFP28	BT/SFP/SFP+/SFP28	BT/SFP/SFP+/SFP28/SFP56	NA
Max Port Speed	50 GbE	100 GbE	100 GbE	NA
NC-SI	Yes	Yes	Yes	NA
SNAPI	No	No	Yes	NA
WoL	Yes	Yes	Yes	NA
Power Consumption	25 W	25 W	15 W - 150 W	25 W for LS
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 both the R6625. For a full list of supported cards and slot priority, please see the PCIe Slot Priority Matrix found in Agile.




Figure 13. Expansion card riser slot connectors

1. Riser 1
2. Riser 2
3. Riser 3
4. Riser 4




Figure 14. Riser 1P

1. Slot 1




Figure 15. Riser 2Q

1. Slot 1
2. Slot 2




Figure 16. Riser 2A

- 1. Slot 1
- 2. Slot 2




Figure 17. Riser 2P

- 1. Slot 1
- 2. Slot 2




Figure 18. Riser 3A

1. Slot 2




Figure 19. Riser 3P

1. Slot 3




Figure 20. Riser 4P

1. Slot 2

Table 17. PCIe Configurations with PERC and rear storage

Config #	RSR Configuration	# of CPUs	PERC type supported	Rear Storage Possible
0	No RSR	2	Front PERC	No
1	R2A + R3A	2	Front PERC	No
2	R2P + R3P	2	Front PERC	No
3	R1P+R4P	2	Front PERC	No
4	R2Q	2	Front PERC	No
5	R3P	2	Front PERC	Yes
7	R2A	1	Front PERC	No

Table 18. PCIe Riser Configurations

Configurations	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width
Config 0	No riser	NA	NA	NA	NA	NA
Config 1. (Riser 2A + Riser 3A)	Riser 2A	1	Processor 1	Low profile	Half length	x16
		2	Processor 2	Low profile	Half length	x16
	Riser 3A	3	Processor 2	Low profile	Half length	x16
Config 2. (Riser 2P + Riser 3P)	Riser 2P	1	Processor 1	Low profile	Half length	x16
		2	Processor 2	Low profile	Half length	x16
	Riser 3P	3	Processor 2	Low profile	Half length	x16
Config 3 (Riser 1P + Riser 4P)	Riser 1P	1	Processor 1	Full Height	Half length	x16
	Riser 4P	2	Processor 2	Full Height	Half length	x16
Config 4 (R2Q)	Riser R2Q	1	Processor 1	Low profile	Half length	x16
		2	Processor 2	Low profile	Half length	x16
Config 5 (Riser R3P)	Riser R3P	3	Processor 2	Low profile	Half length	x16

Table 18. PCIe Riser Configurations (continued)

Configurations	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width
Config 7 (R2A 1 CPU config)	Riser 2A	1	Processor 1	Low profile	Half length	x16

Power, thermal, and acoustics

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby 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](#)
- [Power](#)

Power

Table 19. 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 Enterprise Infrastructure Planning Tool .
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 AMD's GUARDMI 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	AMD's GUARDMI is an embedded technology that provides individual server-level power reporting and power limiting functionality. Dell offers a complete power management solution comprised of AMD's GUARDMI 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.
Fresh Air cooling	See ASHRAE A3/A4 Thermal Restriction.

Table 19. Power tools and technologies (continued)

Feature	Description
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including: • Power distribution units (PDUs) • Uninterruptible power supplies (UPSs) • Energy Smart containment rack enclosures Find additional information at: Data Center Power and Cooling Solutions.

Power

Table 20. 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 EIP Tool .
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.
Fresh Air cooling	Refer to ASHRAE A3/A4 Thermal Restriction.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including: • Power distribution units (PDUs) • Uninterruptible power supplies (UPSs) • Energy Smart containment rack enclosures Find additional information at: PowerEdge servers power and cooling solutions.

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 21. Thermal design characteristics

The thermal design of the PowerEdge R6625 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, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP.
- Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds.
- User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R6625 [Installation and Service Manual](#) and “Advanced Thermal Control: Optimizing across Environments and Power Goals” on Dell.com.
- Cooling redundancy: The PowerEdge R6625 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Cooling redundancy: The PowerEdge R6625 with >4 fans allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the PowerEdge R6625 reliable under a wide range of operating environments.

Acoustics

Acoustical design

Dell PowerEdge delivers sound quality and smooth transient response in addition to sound power levels and sound pressure levels oriented to deployment environments.

Sound quality describes how disturbing or pleasing a person finds a sound, as a function of a variety of psycho-acoustical metrics and thresholds. Tone prominence is one such metric. Transient response refers to how sound changes with time. Sound power level, sound pressure level, and loudness refer to amplitude of sound. A more extensive description of Dell PowerEdge acoustical design and metrics is available in the white paper, [Dell Enterprise Acoustics](#).

Table 21. Acoustical Reference Points and Output Comparisons

Value measured at your ears		Equivalent familiar noise experience
LpA, dBA, re 20μPa	Loudness, sones	
90	80	Loud concert
75	40	Data center, vacuum cleaner, voice must be elevated to be heard
60	10	Conversation levels
45	4	Whispering, open office layout, normal living room
35	2	Quiet office
30	1	Quiet library
20	0	Recording studio

For more information about PowerEdge acoustical design and metrics, see [Understanding Acoustical Data and Causes of Sound in Dell Enterprise Products](#).

Standby Mode

In standby mode, one of the system air movers may continue to run at a low speed in order to provide cooling to components that are powered on when the system is plugged in and you may hear this sound.

Sound Cap

Sound cap is an iDRAC system profile which provides some system performance capping to achieve reduced acoustics without sacrificing reliability. When sound cap is enabled, acoustics are reduced at the expense of system performance. Designed for scenarios in which the server is redeployed from a data center into a noise-sensitive environment, Sound Cap results in limits to acoustical output by applying a power-capping percentage to the CPU(s). The power cap reduces heat generated by the CPUs during high utilization, which reduces fan speed needed for CPU cooling, which in turn reduces acoustical output.

Sound Cap may be useful in the following scenarios:

- Deployment of data center rack servers to quieter environments such as lab or office areas.
- Equipment setup when hardware or software is being loaded.
- Equipment demonstrations or tours in which presenters wish to minimize fan noise from the server.
- When limiting acoustics is consciously prioritized over CPU performance.

Sound Cap applies a percentage-based power cap to the CPU(s) in the server. Sound Cap can limit acoustical output, and it does not impact system cooling performance or any component thermal reliability. It also will not impact fan speeds in moderate CPU workloads or when the system is in idle state. Sound Cap should not be used or may be ineffective in the following scenarios:

- Benchmarking or performance-sensitive applications.
- Reducing idle fan speeds or make a quiet server even quieter.
- Using PCIe-based or VDI workloads.

Sound Cap is enabled in System Settings of the iDRAC GUI or in iDRAC settings within the BIOS setup. Additional information about Sound Cap feature in iDRAC can be found in a white paper published online.

Tips for Acoustically Sensitive Environments

The main cause of computer noise is thermal management, which regulates air mover speeds (such as fans, blowers, etc.) to cool components within their designated limits. Since sound amplitude scales logarithmically with air movement speed, seemingly insignificant changes in speed can have surprisingly noticeable effects. For example, a 10% speed change results in an approximate 2 dB change in sound pressure level, while a 20% speed change results in an approximate 5 dB change in sound pressure level.

In light of this, the table below lists several typical thermal drivers and offers advice on how to deal with them in acoustically sensitive settings. It should be noted that additional components might start to make noise as the ambient temperature, component power, and/or loading drop to the point where air movers are operating at their lowest speed. Examples include the humming of hard drives and the ringing of capacitors.

Table 22. Thermal drivers and tips to address acoustically sensitive environments

Driver	Description	Tip
GPU / FPGA / Accelerator Cards	- GPU, FPGA, or other accelerator cards often push the limits of power and cooling requirements in the PCI form factor. - They may therefore require significantly higher air mover speeds and cause the host system to be much louder. - During system bootup, air mover speeds may ramp (sometimes to full speed) to ensure the cards comply with their thermal limits at assumed worst case thermal state before thermal telemetry feedback is established in bootup. After bootup and telemetry feedback ascertains actual thermal conditions, air mover speeds may reduce.	If deemed an acceptable compromise to accommodate the environment, uninstalling the GPU or running it at a lower loading may reduce air mover speed and acoustical amplitude.
PCI Cards	Dell works diligently with card vendors to validate and develop PCI cards to meet Dell's exacting standards for thermal performance. Although Dell's products can look up and cool appropriately for a variety of cards, some 3 rd party cards may be unknown and thus air mover speed may be higher for thermal protection.	- Replace third party PCI cards with similar Dell supported temperature-controlled cards, if available. - Set target in Third party PCIe card options: Dell provides airflow customization for third-party PCIe adapters installed in PowerEdge platforms. If automatic cooling response is above desired levels (LFM) based on the card specifications, a different LFM target can be set using PCIe Airflow Settings options in iDRAC GUI.
Storage Type	NVME SSD consumes more power than SAS/SATA drive technologies and therefore demands more airflow to achieve system cooling targets thus higher air mover speeds.	If deemed an acceptable compromise to accommodate the environment, replacing NVME SSD devices with SAS/SATA drives may reduce acoustical amplitude.
Thermal Design Power (TDP) System thermal profile selection in BIOS or iDRAC GUI	Higher wattage components under load may require more airflow and thus higher air mover speeds and louder host system. <i>Default Thermal Profile</i>, generally provides a lower air mover speed thus lower acoustical output than those of other thermal profiles. <i>Maximum Performance (Performance Optimized)</i> prioritizes performance above other aspects thus will more quickly increase air mover speed and be louder.	

Table 22. Thermal drivers and tips to address acoustically sensitive environments (continued)

Driver	Description	Tip
	<i>Sound Cap</i>, for products that support the feature and when CPU cooling needs are dictating air mover speeds, can limit maximum acoustical output of the system by sacrificing some CPU performance.	
BOSS-N1 module	If any BOSS-N1 module is installed and "Maximum Performance (Performance Optimized)" is selected, air mover speed and acoustical noise may significantly increase at idle condition	
Ambient Temperature	ISO standards, ECMA standards, and Dell specs require acoustical tests in 23±2°C ambient temperature. When ambient temperature is higher, air movers must speed up to compensate and thus generate more sound.	

PowerEdge R6625 acoustics

Dell PowerEdge R6625 is a rack-mount server for which acoustical output ranges from that appropriate for an office to that of data centers. Acoustical performance is provided in terms of one configuration: Feature Rich. Configuration details are provided in the below table for Feature Rich configuration and acoustical experience data for the configuration is included.

Although the R6625 is intended for data centres, certain users might want a more peaceful environment. In such circumstances, the section on -inch Tips for an Acoustically Sensitive Environment-inch may be useful. However, please be aware that in majority of situations, the system's idle air mover speed cannot be reduced without modifying the system's configuration, and in few instances, even a configuration change may not be able to do so.

Table 23. Feature Rich configuration and acoustical experience data

Configuration	Quietest Low End	Entry	Volume - 1 (HPC)	Feature Rich
CPU Type	AMD Genoa	AMD Genoa	AMD Genoa	AMD Genoa
CPU TDP	200W / 24C	200W / 24C	200W / 24C	290W / 48C
CPU Quantity	2	2	2	2
RDIMM Memory	16 GB DDR5	16 GB DDR5	64 GB DDR5	128 GB DDR5
Memory Quantity	6	6	12	24
Backplane Type	4 x 3.5-inch BP	4 x 3.5-inch BP	10 x 2.5-inch BP	10 x 2.5-inch BP
HDD Type	3.5-inch SATA 2TB	3.5--inch SATA 2TB	2.5-inch NVMe	2.5-inch NVMe
HDD Quantity	2	2	10	10
PSU Type	800 W	800 W	1400 W	1400 W
PSU Quantity	2	2	2	2
M.2	X	X	BOSS-N1	BOSS-N1
OCP	Dual Port 10GbE	Dual Port 10GbE	Dual Port 25GbE	Dual Port 200GbE
PCI 1	X	X	Dual Port 25GbE	X
PCI 2	X	X	Dual Port 25GbE	X
Front PERC	PERC H330	PERC H330	PERC H345	No PERC
LOM Card	X	X	X	X
PERC	X	X	X	X

Table 24. Acoustical performance of R6625 acoustical configurations

Configuration	Quietest Low End	Entry	Volume - 1 (HPC)	Feature Rich
Acoustical Performance: Idle/ Operating @ 25 °C Ambient				

Table 24. Acoustical performance of R6625 acoustical configurations (continued)

Configuration		Quietest Low End	Entry	Volume - 1 (HPC)	Feature Rich
$L_{wA,m}(B)$	Idle	4.7	4.7	5.9	5.9
	Operating	4.9	4.9	5.9	5.9
$K_v(B)$	Idle	0.4	0.4	0.4	0.4
	Operating	0.4	0.4	0.4	0.4
$L_{pA,m}(dB)$	Idle	31	31	43	43
	Operating	32	32	43	43
Prominent tones		No prominent tones in Idle and Operating			
Acoustical Performance: Idle @ 28 °C Ambient					
$L_{wA,m} (B)$		51	51	6.3	6.3
$K_v (B)$		0.4	0.4	0.4	0.4
$L_{pA,m} (dB)$		35	35	47	48
Acoustical Performance: Max. Loading @ 35 °C Ambient					
$L_{wA,m} (B)$		78	78	8.0	8.3
$K_v (B)$		0.4	0.4	0.4	0.4
$L_{pA,m} (dB)$		64	64	63	68

- $L_{wA,m}$: The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 (2017) with data collected using the methods described in ISO 7779 (2010). Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.
- $L_{pA,m}$: 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 described in ISO 7779. The system is placed in a 24U rack enclosure, 25cm above a reflective floor. Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.
- **Prominent discrete tones:** Criteria of Annex D (D.6 and D.11) of ECMA-74 are followed to determine if discrete tones are prominent and to report them, if so.
- **Idle mode:** Idle mode is the steady-state condition in which the server is energized but not operating any intended function.
- **Operating mode:** Operating mode is represented by the maximum of the steady state acoustical output at 50% of CPU TDP or active storage drives for the respective sections of Annex C (C.9.3.2) of ECMA-74.

PowerEdge acoustical dependencies

Some product features impact acoustical server output more than others. The following features are considered strong drivers of acoustical response, thus configurations or operating conditions that include these features may increase air mover speed and acoustical output of the server:

- Ambient temperature: Dell evaluates the acoustical performance of servers in a 23±2°C environment. Ambient temperatures in excess of 25°C will have higher acoustical output and may experience larger fluctuations between state changes.
- Processor thermal design power (TDP): Higher wattage processors may require more airflow to cool under load and thus increase the potential acoustical output of the system.
- Storage type: NVME SSD consumes more power than SAS/SATA drives, and will pre-heat down-stream components (e.g., Processor, DIMM), and therefore demands more airflow to achieve system cooling targets.
- **System thermal profile** selection in BIOS or iDRAC GUI:
 - Default Thermal Profile, generally provides a lower air mover speed thus lower acoustical output than those of other thermal profiles.
 - Maximum Performance (Performance Optimized) will result in higher acoustical output
 - Sound Cap, for products that support the feature, will limit the maximum acoustical output of the system by sacrificing some processor performance.
- BOSS module: If any BOSS module is installed and Maximum Performance (Performance Optimized) is selected, fan speed and acoustical noise may significantly increase at IDLE condition.

Methods to reduce acoustical output of the R6625

Although the R6625 is designed for use in data centers, some users may prefer to use it in a quieter setting. The following is a list of means to do so.

 **NOTE:** Usually, the idle air mover speed of the system cannot be lowered without changing the configuration of the system, and in some cases, even a configuration change may not reduce idle air mover speeds.

- Enable sound cap in iDRAC GUI: Sound cap, a setting in the BIOS, can be toggled on/off during boot up. When enabled, sound cap reduces the acoustics of the system at the expense of some performance.
- Reduce ambient temperature: Lowering the ambient temperature allows the system to cool components more efficiently than at higher ambient temperatures.
- Set target in Third-party PCIe card options: Dell provides airflow customization for third-party PCIe adapters that are installed in PowerEdge platforms. If automatic cooling response is above desired levels (LFM) based on the card specifications, a different LFM target can be set using PCIe Airflow Settings options in iDRAC UI.
- Replace third-party PCI cards with similar Dell supported temperature-controlled cards, if available. Dell works diligently with card vendors to validate and develop PCI cards to meet Dell exacting standards for thermal performance.

Rack, rails, and cable management

Topics:

- [Rails and cable management information](#)

Rails and cable management information

The rail offerings for the PowerEdge R6625 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 [Dell Enterprise Systems Rail Sizing and Rack Compatibility Matrix](#) 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).

A15 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 un-threaded 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 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 interferences with rear mounted PDUs or the rear rack door.




Figure 22. Sliding rails with optional CMA




Figure 23. Sliding rails with optional SRB

A16 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, un-threaded round hole racks including all generations of the Dell racks. Also supports tool-less installation in threaded round hole 4-post racks.
- Support for tool-less installation in Dell Titan or Titan-D racks
- Support full extension of the system out of the rack to allow serviceability of key internal components.
- Support for optional cable management arm (CMA).
- Support for optional strain relief bar (SRB).

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 interferences with rear mounted PDUs or the rear rack door.

A14 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 24. Static rails

Static rails features summary

Static rails for 4-post & 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 un-threaded 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 static rails (A14) must be used. Sliding rails support mounting in 4-post racks only.



Figure 25. 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 (A16) 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 (A14) 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.

i NOTE: CMA is not supported in 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 26. Sliding rails with CMA cabling

Strain Relief Bar (SRB)

The optional strain relief bar (SRB) for the PowerEdge R6625 organizes and supports cable connections at the rear end of the server to avoid damage from bending.




Figure 27. 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 separated 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.




Figure 28. Installing the system in Drop-in sliding rails

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. For a 2U systems, this is a two person lift.

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

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




Figure 29. 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 30. 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 31. 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 32. Pull out the intermediate rail

Table 25. Rail component label

Number	Component
1	Intermediate rail
2	Inner rail

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




Figure 33. Attach the inner rails to the system

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




Figure 34. 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 35. Slide system into the rack

Supported Operating Systems

The PowerEdge system supports the following operating system:

- Canonical Ubuntu Server LTS
- Microsoft Windows Server with Hyper-V
- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server
- VMware vSAN/ESXi
- Citrix XenServer

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 26. 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 27. 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 27. 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 A. Additional specifications

Topics:

- Chassis dimensions
- System weight
- NIC port specifications
- Video specifications
- USB ports specifications
- PSU rating
- Environmental specifications

Chassis dimensions




Figure 36. Chassis dimensions

Table 28. PowerEdge R6625 chassis dimensions

Drives	Xa	Xb	Y	Za with bezel	Za without bezel	Zb	Zc
0 drive	482.0 mm (18.97 inches)	434.0 mm (17.08 inches)	42.8 mm (1.685 inches)	35.84 mm (1.4 inches) With bezel	22.0 mm (0.87 inches) Without bezel	700.7 mm (27.58 inches) Ear to rear wall	736.29 mm (28.99 inches) Ear to PSU handle
4 drives, 8 drives, 10 drives	482.0 mm (18.97 inches)	434.0 mm (17.08 inches)	42.8 mm (1.685 inches)	35.84 mm (1.4 inches) With bezel	22.0 mm (0.87 inches) Without bezel	751.48 mm (29.59 inches) Ear to rear wall	787.05 mm (30.99 inches) Ear to PSU handle

NOTE: Zb is the nominal rear wall external surface where the system board I/O connectors reside.

System weight

Table 29. PowerEdge R6625 system weight

System configuration	Maximum weight (with all drives/SSDs)
A server with fully populated drives	20.4 kg (44.97 pound)
A server without drives and PSU installed	15.8 kg (34.83 pound)

NIC port specifications

The PowerEdge R6625 system supports up to two 10/100/1000 Mbps Network Interface Controller (NIC) ports embedded on the LAN on Motherboard (LOM) and integrated on the Open Compute Project (OCP) cards.

Table 30. NIC port specification for the system

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

NOTE: The system allows either LOM card or an OCP card or both to be installed in the system.

NOTE: On the system board, the supported OCP PCIe width is x8; when x16 PCIe width is installed, it is downgraded to x8.

NOTE: The system allows either LOM card or MIC card to be installed in the system.

Video specifications

The PowerEdge R6625 system supports integrated Matrox G200 graphics controller with 16 MB of video frame buffer.

Table 31. Supported video resolution options

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

USB ports specifications

Table 32. PowerEdge R6625 USB specifications

Front		Rear		Internal (optional)	
USB port type	No. of ports	USB port type	No. of ports	USB port type	No. of ports
USB 2.0-compliant port	One	USB 3.0-compliant port	One	Internal USB 3.0-compliant port	One
iDRAC Direct port (Micro-AB USB 2.0-compliant port)	One	USB 2.0-compliant ports	One		

 **NOTE:** The micro USB 2.0 compliant port can only be used as an iDRAC Direct or a management port.

PSU rating

Table 33. PSUs Highline and Lowline ratings

Features	800 W Platinum	1100 W Titanium	1100 W -48VDC	1400 W Platinum	1400 W Titanium	1800 W Platinum
Peak Power (Highline/-72 VDC)	1360 W	1870 W	1870 W	2380 W	2380 W	3060 W
Highline /-72 VDC	800 W	1100 W	1100 W	1400 W	1400 W	1800 W
Peak Power (Lowline/-40 VDC)	1360 W	1785 W	N/A	1785 W	1785 W	N/A
Lowline /-40 VDC	800 W	1050 W	N/A	1050 W	1050 W	N/A
Highline 240 VDC	800 W	1100 W	N/A	1400 W	1400 W	1800 W
DC -48-60 V	N/A	N/A	1100 W	N/A	N/A	N/A

The PowerEdge R6625 supports up to 2 AC or DC 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 the event that the PSU wattages don't match, the larger of the two PSU's is enabled. Also, there is a PSU mismatch warning displayed in 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 will be flagged as unmatched in iDRAC and the second PSU will not be enabled.

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

Table 34. PSU Efficiency Levels

Efficiency Targets by Load						
Form factor	Output	Class	10%	20%	50%	100%
Redundant 60 mm	800 W AC	Platinum	89.00%	93.00%	94.00%	91.50%
	1100 W AC	Titanium	90.00%	94.00%	96.00%	91.50%
	1100 W -48 VDC	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%

Environmental specifications

Table 35. 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 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 36. 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 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 37. Continuous Operation Specifications for ASHRAE A4

Temperature	Specifications	
Allowable continuous operations		
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 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 38. 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	3,048 meters (10,000 feet)

Table 39. 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 40. 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 41. Particulate contamination specifications

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

Table 42. 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

Particulate and gaseous contamination specifications

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

Table 43. Particulate contamination specifications

Particulate contamination	Specifications
Air filtration: Conventional Data Center only	Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit NOTE: Filtering room air with a MERV8 filter, as specified in ANSI/ASHRAE Standard 127, is a recommended method for achieving the necessary environmental conditions. NOTE: Air entering the data center must have MERV11 or MERV13 filtration. 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.
Walk-Up Edge Data Center or Cabinet (sealed, closed loop environment)	Filtration is not required for cabinets that are anticipated to be opened six 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 43. Particulate contamination specifications (continued)

Particulate contamination	Specifications
Conductive dust: data center and non-data center environments	Air must be free of conductive dust, zinc whiskers, or other conductive particles. i NOTE: Conductive dust, which can interfere with equipment operation, can originate from various sources, including manufacturing processes and zinc whiskers that may develop on the plating of raised floor tiles. i NOTE: This condition applies to data center and non-data center environments.
Corrosive dust: data center and non-data center environments	• Air must be free of corrosive dust. • Residual dust present in the air must have a deliquescent point less than 60% relative humidity. i NOTE: This condition applies to data center and non-data center environments.

Table 44. Gaseous contamination specifications

Gaseous contamination	Specifications	Notes
Copper coupon corrosion rate	ISA-71 Class G1: <300 Å/month	Per ANSI/ISA71.04
Silver coupon corrosion rate	ISA-71 Class G1: <200 Å/month	Per ANSI/ISA71.04

Thermal restriction matrix

Table 45. Processor and heat sink matrix

Heat sink	Notes
T-type HSK	Air-cool HSK

Table 46. Label reference

Label	Description
STD	Standard
HPR Gold	High Performance Gold
HSK	Heat sink
LP	Low profile
FH	Full height

Table 47. Thermal restriction matrix for air cooling configuration

Configuration				4 x 3.5-inch SAS		10 x 2.5-inch SAS		16 E3		14 E3 No fPERC	8 x 2.5-inch NVMe		No BP	Ambient Temperature
Rear configuration				3 LP/ 2 FH	Rear E3	3 LP/ 2 FH	Rear 2 x 2.5-inch SAS Rear E3	3 LP/ 2 FH	2LP	2LP	3 LP/ 2 FH	2LP	3 LP	
CP U TD P	CP U cTD P Max	Model	Core Count	Fan Type										
210 W	240 W	9334	32	STD Fan	STD Fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
200 W	240 W	9254	24	STD Fan	STD Fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
200 W	240 W	9224	24	STD Fan	STD Fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
200 W	240 W	9124	16	STD Fan	STD Fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
290 W	300 W	9634	84	STD Fan*	STD Fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
290 W	300 W	9454	48	STD Fan*	STD Fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
280 W	300 W	9534	64	STD Fan*	STD Fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
280 W	300 W	9354	32	STD Fan*	STD Fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan	HPR Gold fan	HPR Gold fan	35°C
320 W	400 W	9174F	16	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	35°C
320 W	400 W	9274F	24	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	35°C
320 W	400 W	9374F	32	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	35°C
360 W	400 W	9474F	48	N/A	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	35°C
360 W	400 W	9554	64	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	35°C

Table 47. Thermal restriction matrix for air cooling configuration (continued)

Configuration				4 x 3.5-inch SAS		10 x 2.5-inch SAS		16 E3		14 E3 No fPERC	8 x 2.5-inch NVMe		No BP	Ambient Temperature
Rear configuration				3 LP/ 2 FH	Rear E3	3 LP/ 2 FH	Rear 2 x 2.5-inch SAS Rear E3	3 LP/ 2 FH	2LP	2LP	3 LP/ 2 FH	2LP	3 LP	
CP U TD P	CP U cTD P Max	Model	Core Count	Fan Type										
360 W	400 W	9654	96	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	
360 W	400 W	9734	112	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	
360 W	400 W	9754	128	N/A	N/A	N/A	N/A	N/A	N/A	N/A	HPR Gold fan*	HPR Gold fan*	HPR Gold fan*	

NOTE:

- * Supported ambient temperature is 30°C.
- Three dual fan modules are required for single processor, and four dual fan modules are required for dual processor system.
- All air-cooling configurations require CPU shroud.
- OCP Shroud is required to be installed with No Riser Configuration (RC0) when OCP card is present.
- For CPU power > 240 W, DIMM Blank is required to be installed in all empty DIMM slots. For single processor configuration, it follows the same rule for DIMM blanks support.
- For single processor configuration with 3 sets of fan, Fan1 location requires fan blank to be installed if there is no fan.
- For air cooling with 2 x 2.5-inch rear drive configuration, it is required to install Rear Drive Shroud in the system.

Table 48. Thermal restriction matrix for memory configuration

Configuration	4 x 3.5-inch SAS		10 x 2.5-inch SAS		16 E3		14 E3	8 x 2.5-inch NVMe		No BP
Rear configuration	3 LP/ 2 FH	Rear E3	3 LP/ 2 FH	Rear 2 x 2.5-inch SAS	3 LP/ 2 FH	2LP	2LP	3 LP/ 2 FH	2LP	3 LP
DIMM Configuration	STD Fan		HPR Gold fan							
16 GB RDIMM	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
32 GB RDIMM	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
64 GB RDIMM	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C

Table 48. Thermal restriction matrix for memory configuration (continued)

Configuration	4 x 3.5-inch SAS		10 x 2.5-inch SAS		16 E3		14 E3	8 x 2.5-inch NVMe		No BP
Rear configuration	3 LP/ 2 FH	Rear E3	3 LP/ 2 FH	Rear 2 x 2.5-inch SAS	3 LP/ 2 FH	2LP	2LP	3 LP/ 2 FH	2LP	3 LP
DIMM Configuration	STD Fan		HPR Gold fan							
96 GB RDIMM	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
128 GB RDIMM	N/A	N/A	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
256 GB RDIMM	N/A	N/A	30°C	30°C	30°C	30°C	30°C	30°C	30°C	30°C

Thermal restriction for Direct Liquid Cooling

There are no thermal restriction to support all processors for liquid cooling configuration.

 **NOTE:** DIMM blank is not required for liquid cooling configuration. DLC does not support single processor configuration.

Appendix B. Standards compliance

The system conforms to the following industry standards.

Table 49. Industry standard documents

Standard	URL for information and specifications
ACPI Advance Configuration and Power Interface Specification, v6.4	ACPI
Ethernet IEEE Std 802.3-2022	IEEE Standards
MSFT WHQL Microsoft Windows Hardware Quality Labs	Windows Hardware Compatibility Program
IPMI Intelligent Platform Management Interface, v2.0	IPMI
DDR5 Memory DDR5 SDRAM Specification	JEDEC Standards
PCI Express PCI Express Base Specification, v5.0	PCIe Specifications
PMBus Power System Management Protocol Specification, v1.2	Power System Management Protocol Specification
SAS Serial Attached SCSI, 3 (SAS-3) (T10/INCITS 519)	SCSI Storage Interfaces
SATA Serial ATA Rev. 3.3	SATA IO
SMBIOS System Management BIOS Reference Specification, v3.3.0	DMTF SMBIOS
TPM Trusted Platform Module Specification, v1.2 and v2.0	TPM Specifications
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFI Specifications
PI Platform Initialization Specification, v1.7	
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB Implementers Forum, Inc. USB
NVMe Express Base Specification. Revision 2.0c	NVMe
NVMe Command Set Specifications	
1. NVM Express NVM Command Set Specification. Revision 1.1c	
2. NVM Express Zoned Namespaces Command Set. Revision 1.0c	
3. NVM Express® Key Value Command Set. Revision 1.0c	
NVMe Transport Specifications	
1. NVM Express over PCIe Transport. Revision 1.0c	
2. NVM Express RDMA Transport Revision. 1.0b	
3. NVM Express TCP Transport. Revision 1.0c	
NVMe NVM Express Management Interface. Revision 1.2c	
NVMe NVMe Boot Specification. Revision 1.0	

Appendix C Additional resources

Table 50. 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