

# ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU

## Product Guide

The NVIDIA RTX™ 6000 Ada Generation is the ultimate workstation graphics card designed for professionals who demand maximum performance and reliability to deliver their best work and breakthrough innovations across industries. The RTX 6000 provides the unmatched performance and capabilities essential for high-end design, real-time rendering, AI, and high-performance compute workflows.

Built on the NVIDIA Ada Lovelace architecture, the RTX 6000 combines 142 third-generation RT Cores, 568 fourth-generation Tensor Cores, and 18,176 CUDA® cores with 48GB of error correction code (ECC) graphics memory. This all helps deliver the next generation of AI graphics and petaflop inferencing performance for unprecedented speed-up in rendering, AI, graphics, and compute workloads.



Figure 1. ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU

**Did you know?**

NVIDIA RTX professional graphics cards are certified with a broad range of professional applications, tested by leading independent software vendors (ISVs) and workstation manufacturers, and backed by a global team of support specialists. Get the peace of mind to focus on what matters with the premier visual computing solution for mission-critical business.

## Part number information

The following table shows the part numbers for the RTX 6000 Ada GPU.

Table 1. Ordering information

Part number	Feature code	Description	NVIDIA part number	Controlled GPU status
4X67A89324	C2DP	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	900-5G133-2750-001	Controlled

The RTX 6000 Ada GPU is Controlled which means the GPU is not offered in certain markets, as determined by the US Government.

The PCIe option part numbers includes the following:

- One RTX 6000 Ada GPU with full-height (3U) adapter bracket attached
- Documentation

## Features

Key features of the RTX 6000 Ada GPU:

- PCIe Gen 4
- Four DisplayPort 1.4a connectors
- AV1 encode and decode support
- DisplayPort with audio
- 3D stereo support with stereo connector
- NVIDIA GPUDirect® for Video support
- NVIDIA GPUDirect Remote Direct Memory Access (RDMA) support
- NVIDIA virtual GPU (vGPU) software support
- NVIDIA Quadro® Sync II1 compatibility
- NVIDIA RTX Experience
- NVIDIA RTX Desktop Manager software
- NVIDIA RTX IO support
- HDCP 2.2 support
- NVIDIA Mosaic technology

The ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU offers the following features:

- NVIDIA Ada Lovelace Architecture  
NVIDIA RTX™ 6000 Ada Generation is the most powerful workstation GPU offering high-performance, real-time ray tracing, AI-accelerated compute, and professional graphics rendering. Building upon the major SM enhancements from the Ada Lovelace GPU, the NVIDIA Ada Lovelace architecture provides more cores, higher clocks, and a larger L2 cache for more performance to enhance ray tracing operations, tensor matrix operations, and frame rates with DLSS 3.0.
- CUDA Cores  
The NVIDIA Ada Lovelace architecture-based CUDA Cores offer more than 2X the single-precision floating point (FP32) throughput compared to the previous generation, providing significant performance improvements for graphics workflows such as 3D model development and compute for workloads such as desktop simulation for computer-aided engineering (CAE). The RTX 6000 enables two FP32 primary data paths, doubling the peak FP32 operations.
- 3rd Generation RT Cores

Incorporating 3rd generation ray tracing engines, NVIDIA Ada Lovelace architecture-based GPUs provide incredible ray-traced rendering performance. A single RTX 6000 board can render complex professional models with physically accurate shadows, reflections, and refractions to empower users with instant insight. Working in concert with applications leveraging APIs such as NVIDIA OptiX, Microsoft DXR, and Vulkan ray tracing, systems based on the RTX 6000 will power truly interactive design workflows to provide immediate feedback for unprecedented levels of productivity. The RTX 6000 features up to 2X faster ray-triangle intersection throughput compared to the previous generation.

- **4th Generation Tensor Cores**  
Specialized for deep learning matrix multiply and accumulate math operations at the heart of neural network training and inferencing functions, the RTX 6000 includes enhanced Tensor Cores that accelerate more data types and still support the Fine-Grained Structured Sparsity feature that delivers more than 2X throughput for tensor matrix operations compared to the previous generation. New Tensor Cores will accelerate new FP8 precision modes. Independent floating-point and integer data paths allow more efficient execution of workloads using a mix of computation and addressing calculations.
- **5th Generation NVDEC Engine**  
NVDEC is well suited for transcoding and video playback applications for real-time decoding. The following video codecs are supported for hardware-accelerated decoding: MPEG-2, VC-1, H.264 (AVCHD), H.265 (HEVC), VP8, VP9, and AV1 video formats. Video encoding at 8K/60 will be achievable for professional video editing.
- **8th Generation NVENC Engine**  
NVENC can take on the most demanding 4K or 8K video encoding tasks to free up the graphics engine and the CPU for other operations. The RTX 6000 provides better encoding quality than software-based x264 encoders. The RTX 6000 incorporates AV1 video encoding which is 40% more efficient than H.264 encoding for 4K HDR video. AV1 will provide better quality at the same bitrate bandwidth.

## Technical specifications

The following table lists the specifications of the RTX 6000 Ada GPU.

Table 2. RTX 6000 Ada GPU specifications

Feature	Specification
GPU Memory	48 GB GDDR6
Memory Interface	384-bit
Memory Bandwidth	Up to 960 GB/s
ECC	Yes
NVIDIA CUDA Cores	18,176
NVIDIA Tensor Cores	568
NVIDIA RT Cores	142
Single-Precision Performance	91.1 TFLOPS (peak)
RT Core performance	210.6 TFLOPS (peak)
FP8 Tensor Performance	1457.0 TFLOPS (peak, with sparsity)
Host Interface	PCIe 4.0 x16
NVLink support	No
Power Consumption	300 W
Thermal Solution	Active cooling
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height, Full Length
Display Connectors	4x DisplayPort 1.4a
Maximum simultaneous displays	4x 4096 x 2160 @ 120 Hz, 4x 5120 x 2880 @ 60 Hz, 2x 7680 x 4320 @ 60 Hz
Encode / Decode Engines	3x encode, 3x decode (+AV1 encode and decode)
VR Ready	Yes
vGPU software support	NVIDIA vPC/vApps, NVIDIA RTX Virtual Workstation
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6, Vulkan 1.3
Compute APIs	CUDA 11.6, DirectCompute, OpenCL 3.0

## Server support

The following tables list the ThinkSystem servers that are compatible.

Table 3. Server support (Part 1 of 5)

Part Number	Description	AMD V3				2S Intel V3/V4				Multi Node V3	1S V3						
		SR635 V3 (7D9H / 7D9G)	SR655 V3 (7D9F / 7D9E)	SR645 V3 (7D9D / 7D9C)	SR665 V3 (7D9B / 7D9A)	ST650 V3 (7D7B / 7D7A)	SR630 V3 (7D72 / 7D73)	SR650 V3 (7D75 / 7D76)	SR630 V4 (7D68 / 7D69)	SR650 V4 (7D6C / 7D6D)	SR650a V4 (7D6C / 7D6D)	SD535 V3 (7DD8 / 7DD1)	SD530 V3 (7DDA / 7DD3)	SD550 V3 (7DD9 / 7DD2)	ST45 V3 (7DH4 / 7DH5)	ST50 V3 (7DF4 / 7DF3)	ST250 V3 (7DCF / 7DCE)
4X67A89324	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	N	3	N	3	N	N	3	N	N	N	N	N	N	N	N	N

Table 4. Server support (Part 2 of 5)

Part Number	Description	4S 8S Intel V3/V4				GPU Rich				Edge								
		SR650 V3 (7D97 / 7D96)	SR860 V3 (7D94 / 7D93)	SR950 V3 (7DC5 / 7DC4)	SR850 V4 (7DJT / 7DJS)	SR860 V4 (7DJQ / 7DJN)	SR670 V2 (7Z22 / 7Z23)	SR675 V3 (7D9Q / 7D9R)	SR680a V3 (7DHE)	SR680a V3 B200 (7DM9)	SR685a V3 (7DHC)	SR780a V3 (7DJ5)	SR680a V4 (7DMK)	SE100 (7DGR)	SE350 (7Z46 / 7D1X)	SE350 V2 (7DA9)	SE360 V2 (7DAM)	SE450 (7D8T)
4X67A89324	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 5. Server support (Part 3 of 5)

Part Number	Description	Super Computing						1S Intel V2		2S Intel V2		AMD V1						
		SC750 V4 (7DDJ)	SC777 V4 (7DKA)	SD665 V3 (7D9P)	SD665-N V3 (7DAZ)	SD650 V3 (7D7M)	SD650-I V3 (7D7L)	SD650-N V3 (7D7N)	ST50 V2 (7D8K / 7D8J)	ST250 V2 (7D8G / 7D8F)	SR250 V2 (7D7R / 7D7Q)	ST650 V2 (7Z75 / 7Z74)	SR630 V2 (7Z70 / 7Z71)	SR650 V2 (7Z72 / 7Z73)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR645 (7D2Y / 7D2X)	SR665 (7D2W / 7D2V)
4X67A89324	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 6. Server support (Part 4 of 5)

Part Number	Description	Dense V2				4S V2	8S	4S V1		1S Intel V1					
		SD630 V2 (7D1K)	SD650 V2 (7D1M)	SD650-N V2 (7D1N)	SN550 V2 (7Z69)	SR850 V2 (7D31 / 7D32)	SR860 V2 (7Z59 / 7Z60)	SR950 (7X11 / 7X12)	SR850 (7X18 / 7X19)	SR850P (7D2F / 2D2G)	SR860 (7X69 / 7X70)	ST50 (7Y48 / 7Y50)	ST250 (7Y45 / 7Y46)	SR150 (7Y54)	SR250 (7Y52 / 7Y51)
4X67A89324	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 7. Server support (Part 5 of 5)

Part Number	Description	2S Intel V1								Dense V1			
		ST550 (7X09 / 7X10)	SR530 (7X07 / 7X08)	SR550 (7X03 / 7X04)	SR570 (7Y02 / 7Y03)	SR590 (7X98 / 7X99)	SR630 (7X01 / 7X02)	SR650 (7X05 / 7X06)	SR670 (7Y36 / 7Y37)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
4X67A89324	ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU	N	N	N	N	N	N	N	N	N	N	N	N

## Operating system support

The following table lists the supported operating systems:

**Tip:** These tables are automatically generated based on data from [Lenovo ServerProven](#).

Table 8. Operating system support for ThinkSystem NVIDIA RTX 6000 Ada 48GB PCIe Active GPU, 4X67A89324

Operating systems	SR665 V3	SR650 V3 (4th Gen Xeon)	SR655 V3	SR650 V3 (5th Gen Xeon)	SR650 V4/SR650a V4
Microsoft Windows 10	Y	Y	Y	Y	N
Microsoft Windows 11	Y	Y	Y	Y	N
Microsoft Windows Server 2019	Y	Y	Y	Y	N
Microsoft Windows Server 2022	Y	Y	Y	Y	Y
Microsoft Windows Server 2025	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.6	Y	Y	Y	N	N
Red Hat Enterprise Linux 8.7	Y	Y	Y	N	N
Red Hat Enterprise Linux 8.8	Y	Y	Y	Y	N
Red Hat Enterprise Linux 8.9	Y	Y	Y	Y	N
Red Hat Enterprise Linux 8.10	Y	Y	Y	Y	N
Red Hat Enterprise Linux 9.0	Y	Y	Y	N	N
Red Hat Enterprise Linux 9.1	Y	Y	Y	N	N
Red Hat Enterprise Linux 9.2	Y	Y	Y	Y	N
Red Hat Enterprise Linux 9.3	Y	Y	Y	Y	N
Red Hat Enterprise Linux 9.4	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 9.5	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 9.6	Y	Y	Y	Y	N
Red Hat Enterprise Linux 10.0	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 15 SP4	Y	Y	Y	N	N
SUSE Linux Enterprise Server 15 SP4 with Xen	Y	Y	Y	N	N
SUSE Linux Enterprise Server 15 SP5	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 15 SP5 with Xen	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 15 SP6	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 15 SP7	Y	Y	Y	Y	Y
Ubuntu 20.04.5 LTS or later with HWE kernel	Y	N	Y	N	N
Ubuntu 22.04.5 LTS or later with HWE kernel	Y	N	N	N	Y
Ubuntu 22.04 LTS or later with HWE kernel	Y	Y	Y	Y <sup>1</sup>	N
Ubuntu 24.04 LTS or later with HWE kernel	Y	Y	Y	Y	Y

	SR665 V3	SR650 V3 (4th Gen Xeon)	SR655 V3	SR650 V3 (5th Gen Xeon)	SR650 V4/SR650a V4
<b>Operating systems</b>					
VMware vSphere Hypervisor (ESXi) 8.0 U1	Y	Y	Y	N	N
VMware vSphere Hypervisor (ESXi) 8.0 U2	Y	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 8.0 U3	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 9.0	Y	Y	Y	Y	Y

<sup>1</sup> Ubuntu 22.04.3 LTS/Ubuntu 22.04.4 LTS

## NVIDIA GPU software

This section lists the NVIDIA software that is available from Lenovo.

- [NVIDIA vGPU Software \(vApps, vPC, RTX vWS\)](#)
- [NVIDIA Enterprise Software](#)
- [NVIDIA HPC Compiler Software](#)

## NVIDIA vGPU Software (vApps, vPC, RTX vWS)

Lenovo offers the following virtualization software for NVIDIA GPUs:

- **Virtual Applications (vApps)**

For organizations deploying Citrix XenApp, VMware Horizon RDSH or other RDSH solutions. Designed to deliver PC Windows applications at full performance. NVIDIA Virtual Applications allows users to access any Windows application at full performance on any device, anywhere. This edition is suited for users who would like to virtualize applications using XenApp or other RDSH solutions. Windows Server hosted RDSH desktops are also supported by vApps.

- **Virtual PC (vPC)**

This product is ideal for users who want a virtual desktop but need great user experience leveraging PC Windows® applications, browsers and high-definition video. NVIDIA Virtual PC delivers a native experience to users in a virtual environment, allowing them to run all their PC applications at full performance.

- **NVIDIA RTX Virtual Workstation (RTX vWS)**

NVIDIA RTX vWS is the only virtual workstation that supports NVIDIA RTX technology, bringing advanced features like ray tracing, AI-denoising, and Deep Learning Super Sampling (DLSS) to a virtual environment. Supporting the latest generation of NVIDIA GPUs unlocks the best performance possible, so designers and engineers can create their best work faster. IT can virtualize any application from the data center with an experience that is indistinguishable from a physical workstation — enabling workstation performance from any device.

The following license types are offered:

- **Perpetual license**

A non-expiring, permanent software license that can be used on a perpetual basis without the need to renew. For each perpetual license, customers are also required to purchase a 5-year SUMS support contract. Without this contract, the perpetual license cannot be ordered.

- **Annual subscription**

A software license that is active for a fixed period as defined by the terms of the subscription license, typically yearly. The subscription includes Support, Upgrade and Maintenance (SUMS) for the duration of the license term.

- **Concurrent User (CCU)**

A method of counting licenses based on active user VMs. If the VM is active and the NVIDIA vGPU software is running, then this counts as one CCU. A vGPU CCU is independent of the connection to the VM.

The following table lists the ordering part numbers and feature codes.

Table 9. NVIDIA vGPU Software

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
<b>NVIDIA vApps offerings</b>			
7S02005TWW	SEGS	711-VAP001+P3CMI00	vAPPS Perpetual License 1CCU + 5 Years SUMS 1CCU
7S020004WW	B1MQ	711-VAP002+P3CMI12	NVIDIA vApps Subscription License 1 Year, 1 CCU
7S020005WW	B1MR	711-VAP002+P3CMI36	NVIDIA vApps Subscription License 3 Years, 1 CCU
7S02003DWW	S832	711-VAP002+P3CMI48	NVIDIA vApps Subscription License 4 Years, 1 CCU
7S02003EWW	S833	711-VAP002+P3CMI60	NVIDIA vApps Subscription License 5 Years, 1 CCU
<b>NVIDIA vPC offerings</b>			
7S02005UWW	SEGT	711-VPC021+P3CMI00	vPC Perpetual License 1CCU + 5 Years SUMS 1CCU
7S02000AWW	B1MW	711-VPC022+P3CMI12	NVIDIA vPC Subscription License 1 Year, 1 CCU
7S02000BWW	B1MX	711-VPC022+P3CMI36	NVIDIA vPC Subscription License 3 Years, 1 CCU
7S02003FWW	S834	711-VPC022+P3CMI48	NVIDIA vPC Subscription License 4 Years, 1 CCU
7S02003GWW	S835	711-VPC022+P3CMI60	NVIDIA vPC Subscription License 5 Years, 1 CCU
<b>NVIDIA RTX vWS offerings</b>			
7S02005VWW	SEGU	711-DWS021+P3CMI00	vWS Perpetual License 1CCU + 5 Years SUMS 1CCU
7S02005WWW	SEGV	711-DWS021+P3EDI00	vWS EDU Perpetual License 1CCU + 5 Years EDU SUMS 1CCU
7S02000GWW	B1N2	711-DWS022+P3CMI12	NVIDIA RTX vWS Subsc Lic 1Yr 1 CCU
7S02000HWW	B1N3	711-DWS022+P3CMI36	NVIDIA RTX vWS Subscription License 3 Years, 1 CCU
7S02000XWW	S6YJ	711-DWS022+P3CMI48	NVIDIA RTX vWS Subscription License 4 Years, 1 CCU
7S02000YWW	S6YK	711-DWS022+P3CMI60	NVIDIA RTX vWS Subscription License 5 Years, 1 CCU
7S02000MWW	B1N7	711-DWS022+P3EDI12	NVIDIA RTX vWS EDU Subscription License 1 Year, 1 CCU
7S02000NWW	B1N8	711-DWS022+P3EDI36	NVIDIA RTX vWS EDU Subscription License 3 Years, 1 CCU
7S02003BWW	S830	711-DWS022+P3EDI48	NVIDIA RTX vWS EDU Subscription License 4 Years, 1 CCU
7S02003CWW	S831	711-DWS022+P3EDI60	NVIDIA RTX vWS EDU Subscription License 5 Years, 1 CCU
<b>NVIDIA RTX vWS Support &amp; Services</b>			
7S020015WW	S6YS	712-DWSA24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, 1 Year
7S02005CWW	SDZB	712-DWSA24+P3CMI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU 5 Years
7S020016WW	S6YT	712-DWSA24+P3EDI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, EDU, 1 Year
7S02005DWW	SDZC	712-DWSA24+P3EDI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU EDU 5 Years
7S02005EWW	SDZD	712-DWSB24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS SUMS 4 CCU 1 Year
7S020017WW	S6YU	712-DWSD24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Subscription License, 1CCU, 1 Year

## NVIDIA Enterprise Software

Lenovo offers the NVIDIA Enterprise cloud-native enterprise software, comprised of NVIDIA AI Enterprise (NVAIE) and NVIDIA Omniverse Enterprise. NVIDIA Enterprise is an end-to-end, cloud-native suite of AI and data analytics software, optimized, certified, and supported by NVIDIA to run on VMware vSphere and bare-metal with NVIDIA-Certified Systems™. It includes key enabling technologies from NVIDIA for rapid deployment, management, and scaling of AI workloads in the modern hybrid cloud.

NVIDIA Enterprise is licensed on a per-GPU basis. NVIDIA Enterprise products can be purchased as either a perpetual license with support services, or as an annual or multi-year subscription.

- The perpetual license provides the right to use the NVIDIA Enterprise software indefinitely, with no expiration. NVIDIA Enterprise with perpetual licenses must be purchased in conjunction with one-year, three-year, or five-year support services. A one-year support service is also available for renewals.
- The subscription offerings are an affordable option to allow IT departments to better manage the flexibility of license volumes. NVIDIA Enterprise software products with subscription includes support services for the duration of the software's subscription license

The features of NVIDIA Enterprise Software are listed in the following table.

Table 10. Features of NVIDIA Enterprise Software

Features	Supported in NVIDIA Enterprise
Per GPU Licensing	Yes
Compute Virtualization	Supported
Windows Guest OS Support	No support
Linux Guest OS Support	Supported
Maximum Displays	1
Maximum Resolution	4096 x 2160 (4K)
OpenGL and Vulkan	In-situ Graphics only
CUDA and OpenCL Support	Supported
ECC and Page Retirement	Supported
MIG GPU Support	Supported
Multi-vGPU	Supported
NVIDIA GPUDirect	Supported
Peer-to-Peer over NVLink	Supported
GPU Pass Through Support	Supported
Baremetal Support	Supported
AI and Data Science applications and Frameworks	Supported
Cloud Native ready	Supported

Note: Maximum 10 concurrent VMs per product license

The following table lists the ordering part numbers and feature codes.

Table 11. NVIDIA Enterprise Software

Part number	Feature code	Description	NVIDIA part number
	7S02CTO1WW		
AI Enterprise Perpetual License			

<b>Part number</b>	<b>Feature code</b> <b>7S02CTO1WW</b>	<b>Description</b>	<b>NVIDIA part number</b>
7S02001BWW	S6YY	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Perpetual License & Support per GPU, 5 Years	731-AI7004+P3CMI60
7S02001EWW	S6Z1	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Perpetual License & Support per GPU, EDU, 5 Years	731-AI7004+P3EDI60
<b>AI Enterprise Subscription License</b>			
7S02001FWW	S6Z2	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 1 Year	731-AI7003+P3CMI12
7S02005XWW	SENY	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 2 Years	731-AI7003+P3CMI24
7S02001GWW	S6Z3	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 3 Year	731-AI7003+P3CMI36
7S02005YWW	SENZ	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 4 Years	731-AI7003+P3CMI48
7S02001HWW	S6Z4	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 5 Year	731-AI7003+P3CMI60
7S02001JWW	S6Z5	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, EDU, 1 Year	731-AI7003+P3EDI12
7S02005ZWW	SEP0	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, EDU, 2 Years	731-AI7003+P3EDI24
7S02001KWW	S6Z6	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, EDU, 3 Year	731-AI7003+P3EDI36
7S020060WW	SEP1	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, EDU, 4 Years	731-AI7003+P3EDI48
7S02001LWW	S6Z7	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, EDU, 5 Year	731-AI7003+P3EDI60
<b>Business Critical Support Services for NVIDIA Enterprise</b>			
7S02001MWW	S6Z8	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, 1 Year	731-AI7007+P3CMI12
7S02001NWW	S6Z9	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, 3 Year	731-AI7007+P3CMI36
7S020061WW	SEP2	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, 4 Years	731-AI7007+P3CMI48
7S02001PWW	S6ZA	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, 5 Year	731-AI7007+P3CMI60
7S02001QWW	S6ZB	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, EDU, 1 Year	731-AI7007+P3EDI12

Part number	Feature code 7S02CTO1WW	Description	NVIDIA part number
7S02001RWW	S6ZC	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, EDU, 3 Year	731-AI7007+P3EDI36
7S020062WW	SEP3	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, EDU, 4 Years	731-AI7007+P3EDI48
7S02001SWW	S6ZD	Business Critical Support Services for NVIDIA Enterprise (NVIDIA AI Enterprise and Omniverse Enterprise) per GPU, EDU, 5 Year	731-AI7007+P3EDI60

Find more information in the [NVIDIA Enterprise Sizing Guide](#).

## NVIDIA HPC Compiler Software

Table 12. NVIDIA HPC Compiler

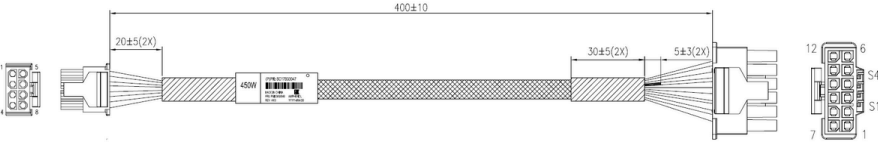
Part number	Feature code 7S09CTO6WW	Description
HPC Compiler Support Services		
7S090014WW	S924	NVIDIA HPC Compiler Support Services, 1 Year
7S090015WW	S925	NVIDIA HPC Compiler Support Services, 3 Years
7S09002GWW	S9UQ	NVIDIA HPC Compiler Support Services, 5 Years
7S090016WW	S926	NVIDIA HPC Compiler Support Services, EDU, 1 Year
7S090017WW	S927	NVIDIA HPC Compiler Support Services, EDU, 3 Years
7S09002HWW	S9UR	NVIDIA HPC Compiler Support Services, EDU, 5 Years
7S090018WW	S928	NVIDIA HPC Compiler Support Services - Additional Contact, 1 Year
7S09002JWW	S9US	NVIDIA HPC Compiler Support Services - Additional Contact, 3 Years
7S09002KWW	S9UT	NVIDIA HPC Compiler Support Services - Additional Contact, 5 Years
7S090019WW	S929	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 1 Year
7S09002LWW	S9UU	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 3 Years
7S09002MWW	S9UV	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 5 Years
HPC Compiler Premier Support Services		
7S09001AWW	S92A	NVIDIA HPC Compiler Premier Support Services, 1 Year
7S09002NWW	S9UW	NVIDIA HPC Compiler Premier Support Services, 3 Years
7S09002PWW	S9UX	NVIDIA HPC Compiler Premier Support Services, 5 Years
7S09001BWW	S92B	NVIDIA HPC Compiler Premier Support Services, EDU, 1 Year
7S09002QWW	S9UY	NVIDIA HPC Compiler Premier Support Services, EDU, 3 Years
7S09002RWW	S9UZ	NVIDIA HPC Compiler Premier Support Services, EDU, 5 Years
7S09001CWW	S92C	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 1 Year
7S09002SWW	S9V0	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 3 Years
7S09002TWW	S9V1	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 5 Years
7S09001DWW	S92D	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 1 Year
7S09002UWW	S9V2	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 3 Years

Part number	Feature code	Description
7S09002VWW	S9V3	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 5 Years

## Auxiliary power cables

The RTX 6000 Ada GPU option part number does not ship with auxiliary power cables. Cables are server-specific due to length requirements. For CTO orders, auxiliary power cables are derived by the configurator. For field upgrades, cables will need to be ordered separately as listed in the table below.

Table 13. Auxiliary power cables for RTX 6000 Ada GPU

Auxiliary power cable needed with the SR650 V3
 <p><b>400mm 16-pin (2x6+4) cable</b>  <b>Feature:</b> BRWK  <b>SBB:</b> SBB7A66338  <b>Base:</b> SC17B33047  <b>FRU:</b> 03KM846  <b>Option:</b>  SR650 V3: 4X67A82883, ThinkSystem SR650 V3 GPU Full Length Thermal Option Kit*</p> <p>* The option part numbers are for thermal kits and include other components needed to install the GPU. See the server product guide for details.</p>

## Regulatory approvals

The RTX 6000 Ada GPU has the following regulatory approvals:

- RCM
- BSMI
- CE
- FCC
- ICES
- KCC
- cUL, UL
- VCCI

## Operating environment

The RTX 6000 Ada GPU has the following operating characteristics:

- Ambient temperature
  - Operational: 0°C to 50°C (-5°C to 55°C for short term\*)
  - Storage: -40°C to 75°C
- Relative humidity:
  - Operational: 5-85% (5-93% short term\*)
  - Storage: 5-95%

\* A period not more than 96 hours consecutive, not to exceed 15 days per year.

## Warranty

One year limited warranty. When installed in a Lenovo server, the GPU assumes the server's base warranty and any warranty upgrades.

## Related publications

For more information, refer to these documents:

- ThinkSystem and ThinkAgile GPU Summary:  
<https://lenovopress.lenovo.com/lp0768-thinksystem-thinkagile-gpu-summary>
- ServerProven compatibility:  
<https://serverproven.lenovo.com/>
- RTX 6000 Ada GPU product page:  
<https://www.nvidia.com/en-us/design-visualization/rtx-6000/>

## Related product families

Product families related to this document are the following:

- [GPU adapters](#)

## Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.  
8001 Development Drive  
Morrisville, NC 27560  
U.S.A.  
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP1940, was created or updated on September 15, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:  
<https://lenovopress.lenovo.com/LP1940>
- Send your comments in an e-mail to:  
[comments@lenovopress.com](mailto:comments@lenovopress.com)

This document is available online at <https://lenovopress.lenovo.com/LP1940>.

## Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ServerProven®

ThinkAgile®

ThinkSystem®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel®, the Intel logo and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, DirectX®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.