

ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU

Product Guide

The NVIDIA RTX A2000 brings the power of NVIDIA RTX technology, realtime ray tracing, AI-accelerated compute, and high-performance graphics to more professionals. Built on the NVIDIA Ampere architecture, the VR ready RTX A2000 combines 26 second-generation RT Cores, 104 third-generation Tensor Cores, and 3,328 next-generation CUDA cores and 6 or 12GB of GDDR6 graphics memory with error correction code (ECC) support for error free computing.

The RTX A2000 GPU features a power-efficient low profile, dual-slot PCIe form factor, and the RTX A2000 12GB doubles memory for even larger models and datasets. Design bigger, render faster, and work smarter than ever before with RTX A2000 GPUs.



Figure 1. ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU

Did you know?

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, NVIDIA RTX is the visual computing solution of choice for demanding enterprise deployments.

Part number information

The following table shows the part numbers for the RTX A2000 GPU.

Table 1. Ordering information

Part number	Feature code	Description	Controlled GPU status
4X67A76720	BMT9	ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU	No

The RTX A2000 GPU is not Controlled which means the GPU is unrestricted and is available in all markets.

The PCIe option part numbers includes the following:

- One RTX A2000 GPU with full-height (3U) adapter bracket attached
- Documentation

Features

The NVIDIA RTX A2000 GPU is a powerful dual slot, low profile GPU solution offering high performance real-time ray tracing, AI-accelerated compute, and professional graphics rendering, all of this in a compact design with unmatched performance. Building upon the major SM enhancements from the Turing GPU, the NVIDIA Ampere architecture enhances ray tracing operations, tensor matrix operations, and concurrent executions of FP32 and INT32 operations.

Key features:

- PCI Express Gen 4
- Four Mini DisplayPort 1.4a connectors with latching mechanism
- AV1 decode support
- DisplayPort with audio
- NVIDIA RTX Experience™
- NVIDIA RTX Desktop Manager software
- HDCP 2.2 support
- NVIDIA Mosaic technology

Technical specifications

The following table lists the specifications of the NVIDIA RTX A2000 GPU.

Table 2. NVIDIA RTX A2000 specifications

Feature	Specification
GPU Memory	12 GB GDDR6
Memory Interface	192-bit
Memory Bandwidth	Up to 288 GB/s
ECC	Yes
NVIDIA CUDA Cores	3,328
NVIDIA Tensor Cores	104
NVIDIA RT Cores	26
Single-Precision Performance	8 TFLOPS (peak)
RT Core performance	15.6 TFLOPS (peak)
Tensor Performance	63.9 TFLOPS (peak)
Host Interface	PCI Express 4.0 x 16
Power Consumption	70 W
Thermal Solution	Active cooling
Form Factor	2.7" H x 6.6" L, Dual Slot, Low Profile
Display Connectors	4x Mini DisplayPort (mDP) 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	1x encode, 2x decode (+AV1 decode)
VR Ready	Yes
Graphics APIs	DirectX 12 Ultimate, 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 4)

Part Number	Description	AMD V3				2S Intel V3/V4				4S 8S Intel V3				Multi Node V3/V4				GPU Rich			
		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 (7DG8 / 7DG9)	SR850 V3 (7D97 / 7D96)	SR860 V3 (7D94 / 7D93)	SR950 V3 (7DC5 / 7DC4)	SD535 V3 (7DD8 / 7DD1)	SD530 V3 (7DDA / 7DD3)	SD550 V3 (7DD9 / 7DD2)	SD520 V4 (7DFZ / 7DFY)	SR670 V2 (7Z22 / 7Z23)	SR675 V3 (7D9Q / 7D9R)	SR680a V3 (7DHE)	SR685a V3 (7DHC)	SR780a V3 (7DJ5)
4X67A76720	ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU	N	3	N	3	4	N	3	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 4. Server support (Part 2 of 4)

Part Number	Description	1S V3			Edge			Super Computing			1S Intel V2		2S Intel V2							
		ST50 V3 (7DF4 / 7DF3)	ST250 V3 (7DCF / 7DCE)	SR250 V3 (7DCM / 7DCL)	SE350 (7Z46 / 7D1X)	SE350 V2 (7DA9)	SE360 V2 (7DAM)	SE450 (7D8T)	SE455 V3 (7DBY)	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)
4X67A76720	ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	4	N	N

Table 5. Server support (Part 3 of 4)

Part Number	Description	AMD V1				Dense V2				4S V2	8S	4S V1		1S Intel V1						
		SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR655 Client OS	SR645 (7D2Y / 7D2X)	SR665 (7D2W / 7D2V)	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)
4X67A76720	ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU	N	N	2	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 6. Server support (Part 4 of 4)

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)
4X67A76720	ThinkSystem NVIDIA RTX A2000 12GB 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 7. Operating system support for ThinkSystem NVIDIA RTX A2000 12GB PCIe Active GPU, 4X67A76720

Operating systems	SR650 V3 (4th Gen Xeon)	SR650 V3 (5th Gen Xeon)	SR655 V3	SR665 V3	ST650 V3 (4th Gen Xeon)	ST650 V3 (5th Gen Xeon)	ST650 V2	SR655
Microsoft Windows 10	Y	Y	Y	Y	N	N	N	N
Microsoft Windows 11	Y	Y	Y	Y	N	N	N	N
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	N
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 7.9	N	N	N	N	N	N	Y	N
Red Hat Enterprise Linux 8.3	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 8.4	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 8.5	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 8.6	Y	N	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 8.7	Y	N	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 8.8	Y	Y	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 8.9	Y	Y	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 9.0	Y	N	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 9.1	Y	N	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 9.2	Y	Y	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 9.3	Y	Y	Y	Y	Y	N	Y	N
SUSE Linux Enterprise Server 15 SP2	N	N	N	N	N	N	Y	N
SUSE Linux Enterprise Server 15 SP3	N	N	N	N	N	N	Y	N
SUSE Linux Enterprise Server 15 SP4	Y	N	Y	Y	Y	N	Y	N
SUSE Linux Enterprise Server 15 SP5	Y	Y	Y	Y	Y	Y	Y	N
Ubuntu 18.04.5 LTS	N	N	N	N	N	N	Y	N
Ubuntu 20.04.5 LTS	N	N	Y	Y	N	N	N	N
Ubuntu 22.04 LTS	Y	Y ¹	Y	Y	Y	N	Y	N
Ubuntu 22.04.3 LTS	N	N	N	N	N	Y	N	N

¹ 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 Omniverse Software \(OVE\)](#)
- [NVIDIA HPC Compiler Software](#)

NVIDIA Omniverse Software (OVE)

NVIDIA Omniverse™ Enterprise is an end-to-end collaboration and simulation platform that fundamentally transforms complex design workflows, creating a more harmonious environment for creative teams.

NVIDIA and Lenovo offer a robust, scalable solution for deploying Omniverse Enterprise, accommodating a wide range of professional needs. This document details the critical components, deployment options, and support available, ensuring an efficient and effective Omniverse experience.

Deployment options cater to varying team sizes and workloads. Using Lenovo NVIDIA-Certified Systems™ and Lenovo OVX nodes which are meticulously designed to manage scale and complexity, ensures optimal performance for Omniverse tasks.

Deployment options include:

- Workstations: NVIDIA-Certified Workstations with RTX 6000 Ada GPUs for desktop environments.
- Data Center Solutions: Deployment with Lenovo OVX nodes or NVIDIA-Certified Servers equipped with L40, L40S or A40 GPUs for centralized, high-capacity needs.

NVIDIA Omniverse Enterprise includes the following components and features:

- Platform Components: Kit, Connect, Nucleus, Simulation, RTX Renderer.
- Foundation Applications: USD Composer, USD Presenter.
- Omniverse Extensions: Connect Sample & SDK.
- Integrated Development Environment (IDE)
- Nucleus Configuration: Workstation, Enterprise Nucleus Server (supports up to 8 editors per scene); Self-Service Public Cloud Hosting using Containers.
- Omniverse Farm: Supports batch workloads up to 8 GPUs.
- Enterprise Services: Authentication (SSO/SSL), Navigator Microservice, Large File Transfer, User Accounts SAML/Account Directory.
- User Interface: Workstation & IT Managed Launcher.
- Support: NVIDIA Enterprise Support.
- Deployment Scenarios: Desktop to Data Center: Workstation deployment for building and designing, with options for physical or virtual desktops. For batch tasks, rendering, and SDG workloads that require headless compute, Lenovo OVX nodes are recommended.

The following part numbers are for a subscription license which is active for a fixed period as noted in the description. The license is for a named user which means the license is for named authorized users who may not re-assign or share the license with any other person.

Table 8. NVIDIA Omniverse Software (OVE)

Part number	Feature 7S02CTO1WW	Description
7S02003ZWW	SCX0	NVIDIA Omniverse Enterprise Subscription per GPU, 1 Year
7S020042WW	SCX3	NVIDIA Omniverse Enterprise Subscription per GPU, 3 Years
7S020044WW	SD5T	NVIDIA Omniverse Enterprise Subscription per GPU, 5 Year
7S020041WW	SCX2	NVIDIA Omniverse Enterprise Subscription per GPU, INC, 1 Year
7S020040WW	SCX1	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 1 Year
7S020043WW	SCX4	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 3 Years
7S020045WW	SD5U	NVIDIA Omniverse Enterprise Subscription per GPU EDU, 5 Year

NVIDIA HPC Compiler Software

Table 9. 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 7S09CTO6WW	Description
7S09002VWW	S9V3	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 5 Years

Auxiliary power cables

The RTX A2000 does not require an auxiliary power cable.

Regulatory approvals

The RTX A2000 GPU has the following regulatory approvals:

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

Operating environment

The RTX A2000 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/>
- NVIDIA RTX A2000 product page:
<https://www.nvidia.com/en-us/design-visualization/rtx-a2000/>
- NVIDIA Ampere Architecture page
<https://www.nvidia.com/en-us/data-center/ampere-architecture/>

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 2024. All rights reserved.

This document, LP1919, 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/LP1919>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1919>.

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® 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.