ThinkSystem and ThinkAgile GPU Summary
Reference Information

Lenovo ThinkSystem servers support GPU technology to accelerate different computing workloads, maximize performance for graphic design, virtualization, artificial intelligence and high performance computing applications in Lenovo servers. This document summarizes the features of the GPUs available for supported ThinkSystem servers and ThinkAgile HX and VX appliances.

Figure 1. ThinkSystem NVIDIA Tesla V100

The following table shows GPUs families and the target workloads

Table 1. GPU families and workloads

Form factor	NVIDIA AI and Virtualization	AMD AI and Virtualization	NVIDIA 3D Graphics
Dual slot	Tesla V100 Tesla P40 Tesla M60 (VDI workloads) Tesla M10 (VDI workloads)	Radeon Pro V340 Instinct MI25	Quadro P6000 Quadro RTX 5000
Single slot	Tesla V100 FHFL Tesla T4		Quadro RTX 4000 Quadro P4000 Quadro P2000 Quadro P620

ThinkSystem server support

The following table summarizes the ThinkSystem server support for the GPUs. The numbers listed in the server columns represent the number of GPUs supported.

Table 2. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	**SE350 (7Z46 / 7D1X)**	**ST50 (7Y48/7Y50)**	**ST250 (7Y45/7Y46)**	**SR150 (7Y54)**	**SR250 (7Y51/7Y52)**	**ST550 (7X09/7X10)**	**SR530 (7X07/7X08)**	**SR550 (7X03/7X04)**	**SR570 (7Y02/7Y03)**	**SR590 (7X98/7X99)**	**SR630 (7X01/7X02)**	**SR650 (7X05/7X06)**	**SR670 (7Y36/7Y37/7Y38)**	**SR635 (7Y98 / 7Y99)**	**SR655 (7Y00 / 7Z01)**	**SR850 (7X18/7X19)**	**SR860 (7X69/7X70)**	**SR950 (7X11/12/13)**	**SD530 (7X21)**	**SD650 (7X58)**	**SN550 (7X16)**	**SN850 (7X15)**
NVIDIA Tesla V100 32GB, 4X67A12088	N	N	N	N	N	N	N	N	N	N	N	2	4	N	3	N	2	N	2	N	N	N
NVIDIA Tesla V100 16GB, 4C57A09498	N	N	N	N	N	N	N	N	N	N	N	2	4	N	3	N	2	N	2	N	N	N
NVIDIA Tesla V100 FHHL, 4X67A11524	N	N	N	N	N	N	N	N	N	N	N	3	N	N	N	N	N	N	N	N	N	N
NVIDIA Tesla P40, 7C57A02888	N	N	N	N	N	N	N	N	N	N	N	2	4	N	N	N	N	N	2	N	N	N
NVIDIA Tesla T4 16GB, 4X67A14926	1	N	N	N	N	N	N	N	N	N	2	5*	8	3	6	N	N	N	N	N	N	N
NVIDIA Tesla M60, 00KG655	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	2	N	2	N	N	N
NVIDIA Tesla M10, 7C57A02891	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	N	N	N
AMD Radeon Pro V340, 4C57A09497	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	N	N	N
AMD Radeon Instinct MI25, 7C57A02897	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	N	N	N
AMD Radeon Instinct MI25 (SR670), 4C57A16224	N	N	N	N	N	N	N	N	N	N	N	N	4	N	N	N	N	N	N	N	N	N
NVIDIA Quadro RTX 5000, 4X67A17267†	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	N	N	N	N
NVIDIA Quadro P6000, 7C57A02895	N	N	N	N	N	2	N	N	N	N	N	2	N	N	N	N	N	N	N	N	N	N
NVIDIA Quadro RTX 4000, 4X67A14934	N	N	N	N	N	2	N	N	N	N	1	2	N	N	N	N	N	N	N	N	N	N
NVIDIA Quadro P4000, 4V17A10255	N	N	N	N	N	2	N	N	N	N	1	3	N	N	N	N	N	N	N	N	N	N
NVIDIA Quadro P2000, 7C57A02877	N	N	1	N	N	2	N	N	N	N	1	N	N	N	N	N	N	N	N	N	N	N
NVIDIA Quadro P620, 4X67A11584	N	1	1	N	1	2	N	N	N	N	3	3	N	N	N	N	N	N	N	N	N	N

* The SR650 has support for 5x T4 or 5x P4 GPUs in servers with second-generation Intel Xeon Scalable processors only. SR650 systems originally with first-generation processors have support for up to 3x T4 or 2x P4 GPUs.

† The NVIDIA Quadro RTX 5000 is only available in countries in EMEA and North America

ThinkAgile HX support

The following tables summarizes the ThinkAgile HX appliance and certified node support for the GPUs. The numbers listed in the server columns represent the number of GPUs supported.

Table 3. ThinkAgile HX appliance and certified node GPU support

Part number	Description	HX Appliances								HX Certified Nodes							
		HX1320 (7X83)	HX1520-R (7X84)	HX3320 (7X83)	HX3520-G (7X84)	HX3720 (7X81)	HX5520 (7X84)	HX5520-C (7X84)	HX7520 (7X84)	HX7820 (7Y95)	HX1321 (7Y89)	HX1521-R (7Y90)	HX3321 (7Y89)	HX3521-G (7Y90)	HX3721 (7Y88)	HX5521 (7Y90)	HX5521-C (7Y90)	HX7521 (7Y90)	HX7821 (7Y96)
4X67A12088	NVIDIA Tesla V100 32GB	N	N	N	2*	N	N	N	N	N	N	N	N	2*	N	N	N	N	N
4C57A09498	NVIDIA Tesla V100 16GB	N	N	N	2*	N	N	N	N	N	N	N	N	2*	N	N	N	N	N
4X67A11524	NVIDIA Tesla V100 FHHL	N	N	N	3*	N	N	N	N	N	N	N	N	3*	N	N	N	N	N
7C57A02888	NVIDIA Tesla P40	N	N	N	2	N	N	N	N	N	N	N	N	2	N	N	N	N	N
4X67A14926	NVIDIA Tesla T4 16GB	N	N	N	5*	N	N	N	N	N	N	N	N	5*	N	N	N	N	N
00KG655	NVIDIA Tesla M60	N	N	N	2	N	N	N	N	N	N	N	N	2	N	N	N	N	N
7C57A02891	NVIDIA Tesla M10	N	N	N	2	N	N	N	N	N	N	N	N	2	N	N	N	N	N
4C57A09497	AMD Radeon Pro V340	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
7C57A02897	AMD Radeon Instinct MI25	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4C57A16224	AMD Radeon Instinct MI25 (SR670)	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4X67A17267	NVIDIA Quadro RTX 5000	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
7C57A02895	NVIDIA Quadro P6000	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4X67A14934	NVIDIA Quadro RTX 4000	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4V17A10255	NVIDIA Quadro P4000	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
7C57A02877	NVIDIA Quadro P2000	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
4X67A11584	NVIDIA Quadro P620	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

* These GPUs are only supported in HX appliances and certified nodes with second-generation Intel Xeon Scalable processors.

ThinkAgile VX support

The following tables summarizes the ThinkAgile VX appliance and certified node support for the GPUs. The numbers listed in the server columns represent the number of GPUs supported.

Table 4. ThinkAgile VX appliance and certified node GPU support

Part number	Description	VX Appliances							Certified Nodes		
		VX2320 (7Y93)	VX3320 (7Y93)	VX3520-G (7Y94)	VX3720 (7Y92)	VX5520 (7Y94)	VX7320-N (7Y93)	VX7520 (7Y94)	VX 1U Cert. Node (7Y93)	VX 2U Cert. Node (7Y94)	VX 2U4N Cert. Node (7Y92)
4X67A12088	NVIDIA Tesla V100 32GB	N	N	N	N	N	N	N	N	2	N
4C57A09498	NVIDIA Tesla V100 16GB	N	N	N	N	N	N	N	N	2	N
4X67A11524	NVIDIA Tesla V100 FHHL	N	N	N	N	N	N	N	N	3	N
7C57A02888	NVIDIA Tesla P40	N	N	N	N	N	N	N	N	2	N
4X67A14926	NVIDIA Tesla T4 16GB	N	N	N	N	N	N	N	N	N	N
7C57A02892	NVIDIA Tesla P4	N	N	N	N	N	N	N	N	N	N
00KG655	NVIDIA Tesla M60	N	N	2	N	N	N	N	N	2	N
7C57A02891	NVIDIA Tesla M10	N	N	2	N	N	N	N	N	2	N
4C57A09497	AMD Radeon Pro V340	N	N	N	N	N	N	N	N	N	N
7C57A02897	AMD Radeon Instinct MI25	N	N	N	N	N	N	N	N	N	N
4C57A16224	AMD Radeon Instinct MI25 (SR670)	N	N	N	N	N	N	N	N	N	N
4X67A17267	NVIDIA Quadro RTX 5000	N	N	N	N	N	N	N	N	N	N
7C57A02895	NVIDIA Quadro P6000	N	N	N	N	N	N	N	N	N	N
4X67A14934	NVIDIA Quadro RTX 4000	N	N	N	N	N	N	N	N	N	N
4V17A10255	NVIDIA Quadro P4000	N	N	N	N	N	N	N	N	N	N
7C57A02877	NVIDIA Quadro P2000	N	N	N	N	N	N	N	N	N	N
4X67A11584	NVIDIA Quadro P620	N	N	N	N	N	N	N	N	N	N

NVIDIA software

Lenovo offers the following software for NVIDIA GPUs:

- GRID 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 GRID 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.

- GRID 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 GRID Virtual PC delivers a native experience to users in a virtual environment, allowing them to run all of their PC applications at full performance.

- Quadro Virtual Datacenter Workstation (vDWS):

  This edition is ideal for mainstream and high-end designers who use powerful 3D content creation applications like Dassault CATIA, SOLIDWORKS, and 3DExcite, Siemens NX, PTC Creo, Schlumberger Petrel, or Autodesk Maya. NVIDIA Quadro Virtual Datacenter Workstation allows users to access their professional graphics applications with full features and performance, anywhere, on 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. Each Lenovo part number includes a fixed number of years of Support, Upgrade and Maintenance (SUMS).

- 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 GRID software is running, then this counts as one CCU. A GRID CCU is independent of the connection to the VM.

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

Table 5. NVIDIA software ordering information

Part number	Feature code*	Description
7S020002WW	B1MN	NVIDIA GRID vApps Perpetual License and SUMS 3Yr, 1 CCU
7S02000PWW	S0DE	NVIDIA GRID vApps Perpetual License and SUMS 4Yr, 1 CCU
7S020003WW	B1MP	NVIDIA GRID vApps Perpetual License and SUMS 5Yr, 1 CCU
7S020004WW	B1MQ	NVIDIA GRID vApps Subscription License 1Yr, 1 CCU
7S020005WW	B1MR	NVIDIA GRID vApps Subscription License 3Yr, 1 CCU
7S020007WW	B1MT	NVIDIA GRID vPC Perpetual License and SUMS 3Yr, 1 CCU
7S020008WW	B1MU	NVIDIA GRID vPC Perpetual License and SUMS 3Yr, 16 CCU
7S02000QWW	S0DF	NVIDIA GRID vPC Perpetual License and SUMS 4Yr, 1 CCU
7S020009WW	B1MV	NVIDIA GRID vPC Perpetual License and SUMS 5Yr, 1 CCU
7S02000AWW	B1MW	NVIDIA GRID vPC Subscription License 1Yr, 1 CCU
7S02000BWW	B1MX	NVIDIA GRID vPC Subscription License 3Yr, 1 CCU
7S02000DWW	B1MZ	NVIDIA Quadro vDWS Perpetual License and SUMS 3Yr, 1 CCU
7S02000EWW	B1N0	NVIDIA Quadro vDWS Perpetual License and SUMS 3Yr, 4 CCU
7S02000RWW	S0DG	NVIDIA Quadro vDWS Perpetual License and SUMS 4Yr, 1 CCU
7S02000FWW	B1N1	NVIDIA Quadro vDWS Perpetual License and SUMS 5Yr, 1 CCU
7S02000GWW	B1N2	NVIDIA Quadro vDWS Subscription License 1Yr, 1 CCU
7S02000HWW	B1N3	NVIDIA Quadro vDWS Subscription License 3Yr, 1 CCU
7S02000KWW	B1N5	NVIDIA Quadro vDWS EDU Perpetual License and SUMS 3Yr, 1 CCU
7S02000SWW	S0DH	NVIDIA Quadro vDWS EDU Perpetual License and SUMS 4Yr, 1 CCU
7S02000LWW	B1N6	NVIDIA Quadro vDWS EDU Perpetual License and SUMS 5Yr, 1 CCU
7S02000MWW	B1N7	NVIDIA Quadro vDWS EDU Subscription License 1Yr, 1 CCU
7S02000NWW	B1N8	NVIDIA Quadro vDWS EDU Subscription License 3Yr, 1 CCU

* For CTO orders, use MTM 7S02CTO1WW

ThinkSystem NVIDIA Tesla V100 GPU

The NVIDIA Tesla V100 GPU Accelerator is an advanced data center GPU ever built to accelerate AI, HPC, and graphics. Powered by NVIDIA Volta, the latest GPU architecture, Tesla V100 offers the performance of up to 100 CPUs in a single GPU, enabling data scientists, researchers, and engineers to tackle challenges that were once thought impossible.

Introduction

The NVIDIA Tesla V100 GPU Accelerator for PCIe is a dual-slot 10.5 inch PCI Express Gen3 card with a single NVIDIA Volta GV100 graphics processing unit (GPU), 16GB or 32GB HBM2 memory with native ECC support. It uses a passive heat sink for cooling. The Tesla V100 GPU supports double precision (FP64), single precision (FP32) and half precision (FP16) compute tasks, unified virtual memory and page migration engine.

Figure 2. ThinkSystem NVIDIA Tesla V100

Part numbers

Table 6. Ordering information

Part number	Feature code	Description
4X67A12088	B34S	ThinkSystem NVIDIA Tesla V100 32GB
4C57A09498	B1JY	ThinkSystem NVIDIA Tesla V100 16GB

Technical specifications

The following table lists the specifications of the NVIDIA Tesla V100 GPU.

Table 7. V100 specifications

Feature	Specification
GPU Architecture	NVIDIA Volta
NVIDIA Tensor Cores	640
NVIDIA CUDA Cores	5,120
Double-Precision Performance	7 TFLOPS
Single-Precision Performance	14 TFLOPS
Tensor Performance	112 TFLOPS
GPU Memory	4X67A12088: 32 GB HBM2 4C57A09498: 16 GB HBM2
Memory Bandwidth	900 GB/sec
ECC	Yes
Interconnect Bandwidth	32 GB/sec
System Interface	PCIe Gen3
Form Factor	PCIe Full Height/Length
Max Power Consumption	250 W
Thermal Solution	Passive
Compute APIs	CUDA, DirectCompute, OpenCL, OpenACC

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 8. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Tesla V100 32GB, 4X67A12088	N	N	N	N	N	N	N	N	N	N	N	2	4	N	3	N	2	N	2	N	N	N
NVIDIA Tesla V100 16GB, 4C57A09498	N	N	N	N	N	N	N	N	N	N	N	2	4	N	3	N	2	N	2	N	N	N

Operating system support

The following tables list the supported operating systems.

Tip: These tables are automatically generated based on data from Lenovo ServerProven.

Table 9. Operating system support for ThinkSystem NVIDIA Tesla V100 32GB GPU, PCIe, 4X67A12088

Operating systems	SR655	SD530 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR860 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)	SR860 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	N	N	Y	Y	N	Y
Microsoft Windows Server 2016	Y	Y	Y	N	Y	Y	Y	N	Y
Microsoft Windows Server 2019	Y	Y	Y	N	Y	Y	Y	N	Y
Microsoft Windows Server version 1803	N	N	N	N	N	N	Y	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	Y	Y	N	Y
Red Hat Enterprise Linux 6.9	N	N	N	N	N	Y	Y	N	Y
Red Hat Enterprise Linux 7.3	N	N	N	N	N	Y	Y	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	Y	Y	N	Y
Red Hat Enterprise Linux 7.5	N	N	N	Y	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	N	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	Y	Y	N	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	Y	Y	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	Y	Y	N	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	N	Y	Y	Y	N	Y
SUSE Linux Enterprise Server 15	N	Y	Y	N	Y	Y	Y	N	Y
SUSE Linux Enterprise Server 15 SP1	Y	N	N	Y	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	Y	Y	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	N	Y	Y	N	Y	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	N	N	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	Y	Y	N	Y	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	N	Y	Y	N	Y	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	N	N	N	N	N	N	N	N

Table 10. Operating system support for ThinkSystem NVIDIA Tesla V100 16GB PCIe Passive GPU, 4C57A09498

Operating systems	SR655	SD530 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR860 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)	SR860 (Gen 1)	x3550 M5 (5463)
Microsoft Windows Server 2012 R2	N	N	N	N	N	Y	Y	N	Y [1]	N
Microsoft Windows Server 2016	Y	Y	Y	N	Y	Y	Y	N	Y [1]	N
Microsoft Windows Server 2019	Y	Y	Y	N	Y	Y	Y	N	Y	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	Y	Y	N	Y	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	Y	Y	N	Y [1]	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	Y	Y	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	Y	Y	N	Y [1]	N
Red Hat Enterprise Linux 7.5	N	N	N	Y	N	Y	Y	Y	Y	N
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	N	Y	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	Y	Y	N	Y [1]	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	Y	Y	N	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	Y	Y	N	Y [1]	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	N	Y	Y	Y	N	Y	N
SUSE Linux Enterprise Server 15	N	Y	Y	N	Y	Y	Y	N	Y	N
SUSE Linux Enterprise Server 15 SP1	Y	N	N	Y	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	Y	Y	N	Y	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	Y	Y	N	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	Y	Y	N	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	N	Y	Y	N	Y	Y	Y	N	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	N	N	N	N	N	N	N	N	N

[1] 2000W PSU x2 and CPU power must less than 200W for NVIDIA V100

ThinkSystem NVIDIA Tesla V100 FHHL GPU

The NVIDIA Tesla V100 FHHL GPU Accelerator is the latest NVIDIA Volta family product, targets for advanced data center to accelerate AI, HPC, and graphics. Tesla V100 FHHL offers significant performance and great power efficiency. The GPU is a full-height half-length (FHHL) form factor.

Introduction

The NVIDIA Tesla V100 FHHL GPU Accelerator for PCIe is a single-slot, full-height half-length (FHHL) PCI Express Gen 3 card with NVIDIA Volta graphics processing unit, 16GB HBM2 memory with native ECC support. It uses a passive heat sink for cooling. The Tesla V100 FHHL PCIe supports double precision (FP64), single precision (FP32) and half precision (FP16) compute tasks, unified virtual memory and page migration engine.

Figure 3. NVIDIA Tesla V100 FHHL GPU

Part numbers

Table 11. Ordering information

Part number	Feature code	Description
4X67A11524	B32D	ThinkSystem NVIDIA Tesla V100 FHHL GPU

Technical specifications

The following table lists the specifications of the NVIDIA Tesla V100 FHHL GPU.

Table 12. V100 FHHL specifications

Feature	Specification
GPU Architecture	NVIDIA Volta
NVIDIA Tensor Cores	640
NVIDIA CUDA® Cores	5,120
Double-Precision Performance	6.5 TFLOPS
Single-Precision Performance	13 TFLOPS
Tensor Performance	105 TFLOPS
GPU Memory	16 GB HBM2
Memory Bandwidth	900 GB/sec
ECC	Yes
Interconnect Bandwidth	32 GB/sec
System Interface	PCIe Gen3
Form Factor	PCIe Full Height/Half Length
Max Power Consumption	150 W
Thermal Solution	Passive
Compute APIs	CUDA, DirectCompute, OpenCL, OpenACC

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 13. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Tesla V100 FHHL, 4X67A11524	N	N	N	N	N	N	N	N	N	N	N	3	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 14. Operating system support for ThinkSystem NVIDIA Tesla V100 FHHL 16GB PCIe Passive GPU , 4X67A11524

Operating systems	SR650 (Gen 2)	SR650 (Gen 1)
Microsoft Windows Server 2012 R2	N	Y
Microsoft Windows Server 2016	Y	Y
Microsoft Windows Server 2019	Y	Y
Red Hat Enterprise Linux 6.10	N	Y
Red Hat Enterprise Linux 6.9	N	Y
Red Hat Enterprise Linux 7.3	N	Y
Red Hat Enterprise Linux 7.4	N	Y
Red Hat Enterprise Linux 7.5	N	Y
Red Hat Enterprise Linux 7.6	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	Y
SUSE Linux Enterprise Server 12 SP2	N	Y
SUSE Linux Enterprise Server 12 SP3	N	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y
SUSE Linux Enterprise Server 15	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y

ThinkSystem NVIDIA Tesla P40 GPU

ThinkSystem NVIDIA Tesla P40 GPU accelerators are purpose-built to deliver maximum throughput for deep learning deployment. Tesla P40 is powered by the revolutionary NVIDIA Pascal architecture provide the computational engine for the new era of artificial intelligence, enabling amazing user experiences by accelerating deep learning applications at scale.

Introduction

ThinkSystem NVIDIA Tesla P40 is a dual-slot 10.5-inch PCIe 3.0 card based on a high-end NVIDIA Pascal graphics processing unit (GPU). It uses a passive heat sink for cooling. P40 has 24GB GDDR5 memory and a 250 W maximum power limit. With 47 TOPS (Tera-Operations Per Second) of inference performance and INT8 operations per GPU, a single server with 8 Tesla P40s delivers the performance of over 140 CPU servers.

Figure 4. ThinkSystem NVIDIA Tesla P40 GPU

Part numbers

Table 15. Ordering information

Part number	Feature code	Description
7C57A02888	AVNZ, B0LZ, B15U*	ThinkSystem NVIDIA Tesla P40

* Feature code varies by server

Technical specifications

The following table lists the specifications of the NVIDIA Tesla P40 GPU.

Table 16. P40 specifications

Feature	Specification
GPU Architecture	NVIDIA Pascal
NVIDIA CUDA® Cores	3840
Single-Precision Performance	12 TeraFLOPS
Integer Operations (INT8)	47 TOPS (Tera-Operations per Second)
GPU Memory	24GB
System Interface	PCIe Gen3
Max Power Consumption	250 W
ECC	Yes
Thermal Solution	Passive
Form Factor	Dual Slot, Full Height
Hardware-Accelerated Video Engine	1x Decode Engine, 2x Encode Engine

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 17. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Tesla P40, 7C57A02888	N	N	N	N	N	N	N	N	N	N	N	2	4	N	N	N	N	N	2	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 18. Operating system support for NVIDIA Tesla P40 GPU, PCIe (passive), 7C57A02888

Operating systems	SD530 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)	nx360 M5 (5465)	x3650 M5 (8871)
Microsoft Windows Server 2012	N	N	N	N	N	N	Y	Y
Microsoft Windows Server 2012 R2	N	N	N	Y	Y	N	Y	Y
Microsoft Windows Server 2016	Y	Y	N	Y	Y	N	Y	Y
Microsoft Windows Server 2019	Y	Y	N	Y	Y	N	N	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 6.10	N	N	N	Y	Y	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	Y	Y	N	N	N
Red Hat Enterprise Linux 7	N	N	N	N	N	N	Y	Y
Red Hat Enterprise Linux 7.3	N	N	N	Y	Y	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y	N	N	N
Red Hat Enterprise Linux 7.5	N	N	Y	Y	Y	Y	N	N
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	N	N	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	Y	Y	N	N	N
SUSE Linux Enterprise Server 11 for AMD64/EM64T	N	N	N	N	N	N	Y	Y
SUSE Linux Enterprise Server 12	N	N	N	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y	N	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	Y	Y	N	N	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	N	Y	Y	N	N	N
SUSE Linux Enterprise Server 15	Y	Y	N	Y	Y	N	N	Y
SUSE Linux Enterprise Server 15 SP1	N	N	Y	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	Y	Y	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	N	Y	Y	N	N	N

ThinkSystem NVIDIA Tesla T4 GPU

The NVIDIA Tesla T4 GPU supports diverse cloud workloads, including high-performance computing, deep learning training and inference, machine learning, data analytics, and graphics. Based on the new NVIDIA Turing Architecture and packaged in an energy-efficient 70-watt, small PCIe form factor, Tesla T4 is optimized for scale-out computing environments with its multi-precision Turing Tensor Cores and new RT Cores.

Introduction

The NVIDIA Tesla T4 GPU is a single-slot, low-profile, 6.6-inch PCIe 3.0 Universal Deep Learning Accelerator. The Tesla T4 has 16GB GDDR6 memory.

Tesla T4 provides revolutionary multi-precision performance to accelerate deep learning and machine learning training and inference, video transcoding and virtual desktops. Powering performance from FP32 to FP16 to INT8, as well as INT4 precisions, T4 delivers up to 40x higher performance than CPUs. As part of the NVIDIA AI platform, Tesla T4 supports all AI frameworks and network types, delivering dramatic performance and efficiency.

Figure 5. ThinkSystem NVIDIA Tesla T4 GPU

Part numbers

Table 19. Ordering information

Part number	Feature code	Description
4X67A14926	B4YB	ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU

Technical specifications

The following table lists the specifications of the NVIDIA Tesla T4 GPU.

Table 20. Technical Specifications

Feature	Specification
GPU Architecture	NVIDIA Turing
NVIDIA Turing Tensor Cores	320 cores
NVIDIA CUDA Cores	2560 cores
Single Precision Performance (FP32)	8.1 TFLOPS
Mixed Precision (FP16/FP32)	65 FP16 TFLOPS
INT8 Precision	130 INT8 TOPS
INT4 Precision	260 INT4 TOPS
System Interface	PCIe 3.0 x16 or x8 (x16 physical connector)
GPU Memory	16 GB GDDR6
Memory Bandwidth	320+ GB/s
Form Factor	single-slot, low-profile
Max Power Consumption	70 W
Thermal Solution	Passive
Compute APIs	CUDA, DirectCompute, OpenCL, OpenACC

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 21. ThinkSystem server support

	E	1S Intel				2S Intel							AMD		4S Intel			Dense/ Blade			
Description & part number	**SE350 (7Z46 / 7D1X)**	**ST50 (7Y48/7Y50)**	**ST250 (7Y45/7Y46)**	**SR150 (7Y54)**	**SR250 (7Y51/7Y52)**	**ST550 (7X09/7X10)**	**SR530 (7X07/7X08)**	**SR550 (7X03/7X04)**	**SR570 (7Y02/7Y03)**	**SR590 (7X98/7X99)**	**SR630 (7X01/7X02)**	**SR650 (7X05/7X06)**	**SR670 (7Y36/7Y37/7Y38)**	**SR635 (7Y98 / 7Y99)**	**SR655 (7Y00 / 7Z01)**	**SR850 (7X18/7X19)**	**SR860 (7X69/7X70)**	**SR950 (7X11/12/13)**	**SD530 (7X21)**	**SD650 (7X58)**	**SN550 (7X16)**	**SN850 (7X15)**
NVIDIA Tesla T4 16GB, 4X67A14926	1	N	N	N	N	N	N	N	N	N	2	5*	8	3	6	N	N	N	N	N	N	N

* The SR650 has support for 5x T4 GPUs in servers with second-generation Intel Xeon Scalable processors only. SR650 systems originally with first-generation processors have support for up to 3x T4 GPUs.

Operating system support

The following table lists the supported operating systems.

Tip: These tables are automatically generated based on data from Lenovo ServerProven.

Table 22. Operating system support for ThinkSystem NVIDIA Tesla T4 75W PCIe 16GB, 4X67A14926

Operating systems	SE350	SR635	SR655	SR630 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR630 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	Y	Y	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	N	Y	Y	N
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	N	Y	Y	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	Y	Y	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	Y	Y	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	Y	Y	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	Y	Y	N
Red Hat Enterprise Linux 7.5	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	Y	Y	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	Y	Y	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	N	Y	Y	N
SUSE Linux Enterprise Server 12 SP4	N	Y	Y	Y	Y	N	Y	Y	N
SUSE Linux Enterprise Server 15	Y	N	N	Y	Y	N	Y	Y	N
SUSE Linux Enterprise Server 15 SP1	N	Y	Y	N	N	Y	Y	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	N	N	Y	Y	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U3	N	Y	Y	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	N	N	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	N	N	Y	Y	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	N	N	Y	Y	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.7 U3	N	Y	Y	N	N	N	N	N	N

ThinkSystem NVIDIA Tesla M60 GPU

ThinkSystem NVIDIA Tesla M60 GPU accelerator works with NVIDIA GRID software to provide the industry's highest user performance for virtualized applications. This solution allows enterprises to virtualize any application— including professional graphics applications—and deliver them out to any device, anywhere.

Introduction

ThinkSystem NVIDIA Tesla M60 is a dual-slot 10.5 inch PCI Express Gen3 graphics card with two high-end NVIDIA Maxwell graphics processing units (GPUs). The Tesla M60 has 16 GB GDDR5 memory (8 GB per GPU) and a 300 W maximum power limit. M60 is offered as a 300 W passively cooled adapter that uses system airflow to keep the card within its thermal limits.

Figure 6. ThinkSystem NVIDIA Tesla M60 GPU

Part numbers

Table 23. Ordering information

Part number	Feature code	Description
00KG655	ATZE, B0TB, B13J, B0M0*	ThinkSystem NVIDIA Tesla M60

* Feature code varies by server

Technical specifications

The following table lists the specifications of the NVIDIA Tesla M60 GPU.

Table 24. M60 specifications

Feature	Specification
Virtualization Use Case	Performance-Optimized, Graphics Virtualization
GPU Architecture	NVIDIA Maxwell
GPUs per Board	2
Max User per Board	32 (16 per GPU)
NVIDIA CUDA® Cores	4096 NVIDIA CUDA Cores (2048 per GPU)
GPU Memory	16 GB of GDDR5 Memory (8 per GPU)
H.264 1080p30 streams	36
Max Power Consumption	300 W
Thermal Solution Active/Passive	Passive
Form Factor PCIe 3.0 Dual Slot	PCIe 3.0 Dual Slot

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Server memory limitation: When the M60 is installed, the total system memory must be less than 1 TB. See https://support.lenovo.com/us/en/solutions/ht114952 for details.

Table 25. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Tesla M60, 00KG655	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	2	N	2	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 26. Operating system support for NVIDIA Tesla M60 GPU, PCIe (Passive), 00KG655

Operating systems	SD530 (Gen 2)	SR650 (Gen 2)	SR860 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)	SR860 (Gen 1)	nx360 M5 (5465)
Microsoft Windows Server 2012	N	N	N	N	N	N	Y
Microsoft Windows Server 2012 R2	N	N	N	Y	Y	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	Y
Red Hat Enterprise Linux 6.10	N	N	N	Y	Y	Y	N
Red Hat Enterprise Linux 6.9	N	N	N	Y	Y	Y	N
Red Hat Enterprise Linux 7	N	N	N	N	N	N	Y
Red Hat Enterprise Linux 7.3	N	N	N	Y	Y	N	N
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y	Y	N
Red Hat Enterprise Linux 7.5	N	N	N	Y	Y	Y	N
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	Y	Y	Y	N
SUSE Linux Enterprise Server 11 for AMD64/EM64T	N	N	N	N	N	N	Y
SUSE Linux Enterprise Server 12	N	N	N	N	N	N	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	Y	Y	Y	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	N
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	Y
VMware vSphere 5.1 (ESXi)	N	N	N	N	N	N	Y
VMware vSphere Hypervisor (ESXi) 5.5	N	N	N	N	N	N	Y
VMware vSphere Hypervisor (ESXi) 6.0	N	N	N	N	N	N	Y
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	Y	Y	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	N

ThinkSystem NVIDIA Tesla M10 GPU

ThinkSystem NVIDIA Tesla M10 GPU accelerator works with NVIDIA GRID software to provide the industry's highest user density for virtualized desktops and applications. It supports 64 desktops per board and 128 desktops per server, giving your business the power to deliver great experiences to all of your employees at an affordable cost.

Introduction

ThinkSystem NVIDIA Tesla M10 is a dual-slot 10.5 inch PCI Express 3.0 graphics card with four mid-range NVIDIA Maxwell graphics processing units (GPUs). The Tesla M10 has 32 GB GDDR5 memory (8 GB per GPU) and a 225 W maximum power limit. The board is passively cooled and supports both airflow directions.

Figure 7. ThinkSystem NVIDIA Tesla M10 GPU

Part numbers

Table 27. Ordering information

Part number	Feature code	Description
7C57A02891	AX8L, B15V, B0RK*	ThinkSystem NVIDIA Tesla M10

* Feature code varies by system

Technical specifications

The following table lists the specifications of the NVIDIA Tesla M10 GPU.

Table 28. M10 specifications

Feature	Specification
Virtualization Use Case	Performance-Optimized Graphics Virtualization
GPU Architecture	NVIDIA Maxwell
GPUs per Board	4
Max User per Board	64 (16 per GPU)
NVIDIA CUDA® Cores	2560 NVIDIA CUDA Cores (640 per GPU)
GPU Memory	32 GB of GDDR5 Memory (8 per GPU)
H.264 1080p30 streams	28
Max Power Consumption	225 W
Thermal Solution Active/Passive	Passive
Form Factor PCIe 3.0 Dual Slot	PCIe 3.0 Dual Slot

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Server memory limitation: When the M10 is installed, the total system memory must be less than 1 TB. See https://support.lenovo.com/us/en/solutions/ht114952 for details.

Table 29. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Tesla M10, 7C57A02891	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	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 30. Operating system support for NVIDIA Tesla M10 32GB GPU, PCIe (Passive), 7C57A02891

Operating systems	SD530 (Gen 2)	SR650 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)	nx360 M5 (5465)	x3650 M5 (8871)	x240 M5 (9532)
Microsoft Windows Server 2012	N	N	N	N	Y	Y	N
Microsoft Windows Server 2012 R2	N	N	N	N	Y	Y	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	N
Microsoft Windows Server 2019	Y	Y	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0	N	N	N	N	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	Y	Y	Y	Y	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7	N	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	N	N	N

ThinkSystem AMD Radeon Pro V340 GPU

The AMD Radeon Pro V340 datacenter graphics card delivers an impressively smooth GPU experience from the cloud to virtually any device, anywhere. With MxGPU technology at its core, this hardware-based virtualized graphics solution provides high levels of predictable performance, enhanced security and support for up to 32 VMs per card. This MxGPU solution is easy to set up and manage, and does not require end user licenses, providing enterprises with a lower cost per user.

Introduction

The AMD Radeon Pro V340 delivers a smooth GPU experience to every virtual application and desktop. It is a dual GPU card that is purpose-built to deliver extreme performance along with maximum user density for virtualized graphics, delivered from on-premise VDI environments or from the cloud. Based on AMD's MxGPU (Multiuser) technology, the Radeon Pro V340 gives every VM the benefits of having a high performance graphics card, giving users the visual experience they expect without having an actual GPU in their system.

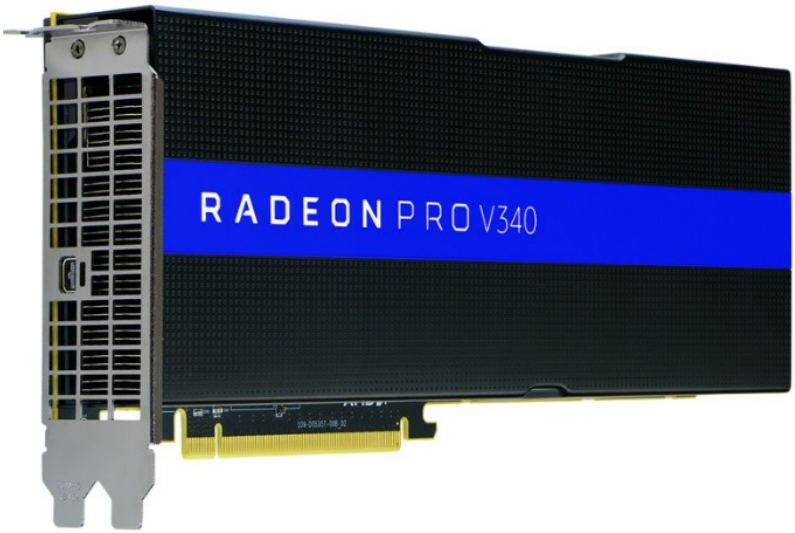

Figure 8. ThinkSystem AMD Radeon Pro V340 GPU

Part numbers

Table 31. Ordering information

Part number	Feature code	Description
4C57A09497	B32P	ThinkSystem AMD Radeon Pro V340 32GB PCIe Passive GPU

Technical specifications

The following table lists the specifications of the AMD Radeon Pro V340 GPU.

Table 32. Technical specifications

Feature	Specification
GPU Architecture	AMD Vega
Lithography	14nm FinFET
Stream Processors	7168
Compute Units	112
Virtual Machines per card	32
System Interface	PCIe 3.0 x16
GPU Memory	32 GB HBM2
Memory Bandwidth	484 GB/s
Form Factor	Double Slot, Full Height
Typical Board Power	<300 W
Thermal Solution	Passive

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 33. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
AMD Radeon Pro V340, 4C57A09497	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	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 34. Operating system support for ThinkSystem AMD Radeon Pro V340 32GB PCIe Passive GPU, 4C57A09497

Operating systems	SR650 (Gen 2)	SR650 (Gen 1)
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y

ThinkSystem AMD Radeon Instinct MI25 GPU

The Radeon Instinct MI25 accelerator, designed with the most advanced Next-Gen “Vega” GPU architecture, is the ultimate training accelerator for large scale machine intelligence and deep learning, along with being an optimized open compute workhorse for single-precision HPC-class system workloads.

Introduction

The MI25 delivers leadership in FP16 and FP32 performance in a passively-cooled single GPU server card with 24.6 TFLOPS of FP16 and 12.3 TFLOPS of FP32 peak performance through its 64 compute units with 4,096 stream processors.

The Radeon Instinct MI25’s powerful compute engine and advanced memory architecture, combined with AMD’s ROCm open software platform and ecosystem, provides a powerful, flexible heterogeneous compute solution that allows datacenter designers to meet the challenges of a new era of compute and Machine Intelligence.

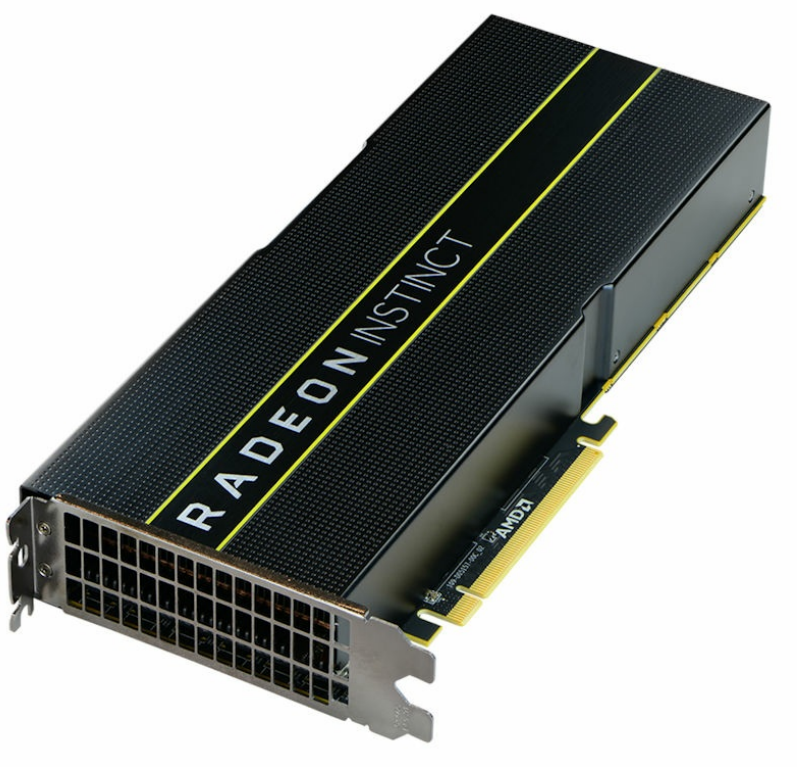

Figure 9. AMD Radeon Instinct MI25 GPU

Part numbers

Table 35. Ordering information

Part number	Feature code	Description
7C57A02897	B228	ThinkSystem AMD Radeon Instinct MI25 16GB PCIe Passive GPU
4C57A16224	B5DN	ThinkSystem SR670 AMD Radeon Instinct MI25 16Gb PCIe Passive GPU

Technical specifications

The following table lists the specifications of the AMD Radeon Instinct MI25 GPU.

Table 36. Technical specifications

Feature	Specification
GPU Architecture:	AMD “Vega10”
Stream Processors:	4096
GPU Memory:	16GB HBM2
Memory Bandwidth:	Up to 484 GB/s
Half-Precision Performance (FP16)	24.6 TFLOPS
Single-Precision Performance (FP32)	12.3 TFLOPS
Double-Precision Performance (FP64)	768 GFLOPS
ECC:	Yes
Bus Interface:	PCIe Gen 3 x16
MxGPU Capability:	Yes
Board Form Factor:	Full-Height, Dual-Slot
Thermal Solution:	Passively Cooled
Standard Max Power:	300W TDP
OS Support:	Linux 64-bit
ROCm Software Platform:	Yes
Programing Environment:	ISO C++, OpenCL, CUDA (via AMD’s HIP conversion tool) and Python5 (via Anaconda’s NUMBA)

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 37. ThinkSystem server support

	E	1S Intel				2S Intel							AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
AMD Radeon Instinct MI25, 7C57A02897	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	2	N	N	N
AMD Radeon Instinct MI25 (SR670), 4C57A16224	N	N	N	N	N	N	N	N	N	N	N	N	4	N	N	N	N	N	N	N	N	N

Operating system support

The following tables list the supported operating systems.

Tip: These tables are automatically generated based on data from Lenovo ServerProven.

Table 38. Operating system support for ThinkSystem AMD Radeon Instinct MI25 16GB PCIe Passive GPU, 7C57A02897

Operating systems	SD530 (Gen 2)	SR650 (Gen 2)	SD530 (Gen 1)	SR650 (Gen 1)
Red Hat Enterprise Linux 7.3	N	N	Y	Y
Red Hat Enterprise Linux 7.4	N	N	Y	Y
Red Hat Enterprise Linux 7.5	N	N	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y

Table 39. Operating system support for ThinkSystem SR670 AMD Radeon Instinct MI25 16Gb PCIe Passive GPU, 4C57A16224

Operating systems	SR670 (Gen 2)	SR670 (Gen 1)
Red Hat Enterprise Linux 7.5	Y	Y
Red Hat Enterprise Linux 7.6	Y	N

ThinkSystem NVIDIA Quadro RTX 5000 GPU

Shatter the boundaries of what's possible with the NVIDIA Quadro RTX 5000, powered by NVIDIA Turing GPU to bring real-time ray tracing and accelerated AI to next-generation workflows. Creative and technical professionals can supercharge demanding design and visualization workloads and make more informed decisions faster than ever before.

EMEA and NA only: The NVIDIA Quadro RTX 5000 is only available in countries in EMEA and North America.

Introduction

The Quadro RTX 5000 is equipped with 3072 CUDA cores, 384 Tensor cores, 48 RT cores and 16GB GDDR6 memory. It can render complex models and scenes with physically accurate shadows, reflections, and refractions to empower users with instant insight.

Figure 10. ThinkSystem NVIDIA Quadro RTX 5000 GPU

Part numbers

Table 40. Ordering information

Part number	Feature code	Description
4X67A17267	B6CH	ThinkSystem NVIDIA Quadro RTX 5000 16GB PCIe Active GPU

Technical specifications

The following table lists the specifications of the NVIDIA Quadro RTX 5000 GPU.

Table 41. RTX 5000 specifications

Feature	Specification
GPU Memory	16 GB GDDR6
Memory Interface	256-bit
Memory Bandwidth	Up to 448 GB/s
NVIDIA CUDA® Cores	3072
NVIDIA Tensor Cores	384
NVIDIA RT Cores	48
System Interface	PCI Express 3.0 x16
Max Power Consumption	265 W
Thermal Solution	Active
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height
Display Connectors	4x DP 1.4, 1x USB-C
Max Simultaneous Displays	4 direct, 4 DP1.4 Multi-Stream
Max DP 1.4 Resolution	4x 4096x2160 @ 120 Hz 4x 5120x2880 @ 60 Hz 2x 7680x4320 @ 60 Hz
Graphics APIs	DirectX 12.0, Shader Model 5.1, OpenGL 4.6, Vulkan 1.1
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 42. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro RTX 5000, 4X67A17267†	N	N	N	N	N	N	N	N	N	N	N	2	N	N	N	N	N	N	N	N	N	N

† The NVIDIA Quadro RTX 5000 is only available in countries in EMEA and North America

Operating system support

The following table lists the supported operating systems.

Tip: This table is automatically generated based on data from Lenovo ServerProven.

Table 43. Operating system support for ThinkSystem NVIDIA Quadro RTX 5000 16GB PCIe Active GPU, 4X67A17267

Operating systems	SR650 (Gen 2)	SR650 (Gen 1)
Microsoft Windows Server 2012 R2	N	Y
Microsoft Windows Server 2016	Y	Y
Microsoft Windows Server 2019	Y	Y
Red Hat Enterprise Linux 7.4	N	Y
Red Hat Enterprise Linux 7.5	N	Y
Red Hat Enterprise Linux 7.6	Y	Y
Red Hat Enterprise Linux 8.0	N	Y
SUSE Linux Enterprise Server 12 SP2	N	Y
SUSE Linux Enterprise Server 12 SP3	N	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y
SUSE Linux Enterprise Server 15	Y	Y
SUSE Linux Enterprise Server 15 SP1	N	Y

ThinkSystem NVIDIA Quadro P6000 GPU

ThinkSystem NVIDIA Quadro P6000 GPU cards are certified with a broad range of sophisticated professional applications, tested by leading workstation manufacturers, and backed by a global team of support specialists. This gives you the peace of mind to focus on doing your best work. Whether you're developing revolutionary products or telling spectacularly vivid visual stories, Quadro gives you the performance to do it brilliantly.

Introduction

ThinkSystem NVIDIA Quadro P6000 is the world's most advanced professional graphics solution ever created, combining the latest GPU, memory and display technologies that result in unprecedented performance and breakthrough capabilities. Professionals across of range of industries can now create their most complex designs, solve the most challenging visualization problems and experience their creations within the most detailed, life-like VR environments.

Figure 11. ThinkSystem NVIDIA Quadro P6000 GPU

Part numbers

Table 44. Ordering information

Part number	Feature code	Description
7C57A02895	AVRK, B2V5*	ThinkSystem NVIDIA Quadro P6000

* Feature code varies by system

Technical specifications

The following table lists the specifications of the NVIDIA Quadro P6000 GPU.

Table 45. P6000 specifications

Feature	Specification
GPU Memory	24 GB GDDR5X
Memory Interface	384-bit
Memory Bandwidth	Up to 432 GB/s
NVIDIA CUDA® Cores	3840
System Interface	PCI Express 3.0 x16
Max Power Consumption	250 W
Thermal Solution	Active
Form Factor	4.4''H x 10.5'' L, Dual Slot, Full Height
Display Connectors	4x DP 1.4 + DVI-D DL
Max Simultaneous Displays	4 direct, 4 DP1.4 Multi-Stream
Max DP 1.4 Resolution	7680 x 4320 @ 30 Hz
Max DVI-D DL Resolution	2560 x 1600 @ 60 Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.54, DirectX 12.05, Vulkan 1.04
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 46. ThinkSystem server support

	E	1S Intel				2S Intel							AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro P6000, 7C57A02895	N	N	N	N	N	2	N	N	N	N	N	2	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 47. Operating system support for ThinkSystem NVIDIA Quadro P6000 24GB PCIe Active GPU, 7C57A02895

Operating systems	SR650 (Gen 2)	ST550 (Gen 2)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y
Red Hat Enterprise Linux 6.10	N	N	Y	Y
Red Hat Enterprise Linux 6.9	N	N	Y	Y
Red Hat Enterprise Linux 7.3	N	N	Y	Y
Red Hat Enterprise Linux 7.4	N	N	Y	Y
Red Hat Enterprise Linux 7.5	N	N	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	N	Y	N	N
SUSE Linux Enterprise Server 11 SP4	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y

ThinkSystem NVIDIA Quadro RTX 4000 GPU

Meet the challenge of today's demanding professional workflows with NVIDIA Quadro RTX 4000, powered by NVIDIA Turing architecture and the NVIDIA RTX platform. The NVIDIA Quadro RTX 4000 delivers GPU accelerated ray tracing, deep learning, and advanced shading in an accessible single slot form factor. It gives designers the power to accelerate their creative efforts with faster time to insight and faster time to solution.

Introduction

The Quadro RTX 4000 is equipped with 2304 CUDA cores, 288 Tensor Cores, 36 RT cores and 8 GB GDDR6 memory. The single-slot GPU is designed to manage the most intensive AEC, DCC, AI, VR and graphics workloads.

Figure 12. ThinkSystem NVIDIA Quadro RTX 4000 GPU

Part numbers

Table 48. Ordering information

Part number	Feature code	Description
4X67A14934	B6CG	ThinkSystem NVIDIA Quadro RTX 4000 8GB PCIe Active GPU

Technical specifications

The following table lists the specifications of the NVIDIA Quadro RTX 4000 GPU.

Table 49. RTX 4000 specifications

Feature	Specification
GPU Memory	8 GB GDDR6
Memory Interface	256-bit
Memory Bandwidth	Up to 416 GB/s
NVIDIA CUDA® Cores	2304
NVIDIA Tensor Cores	288
NVIDIA RT Cores	36
System Interface	PCI Express 3.0 x16
Max Power Consumption	160 W
Thermal Solution	Active
Form Factor	4.4" H x 9.5" L, Single Slot, Full Height
Display Connectors	3x DP 1.4, 1x USB-C
Max Simultaneous Displays	4 direct, 4 DP1.4 Multi-Stream
Max DP 1.4 Resolution	4x 3840x2160 @ 120 Hz 4x 5120x2880 @ 60 Hz 2x 7680x4320 @ 60 Hz
Graphics APIs	DirectX 12.0, Shader Model 5.1, OpenGL 4.6, Vulkan 1.1
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 50. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro RTX 4000, 4X67A14934	N	N	N	N	N	2	N	N	N	N	1	2	N	N	N	N	N	N	N	N	N	N

Operating system support

The following table lists the supported operating systems.

Tip: This table is automatically generated based on data from Lenovo ServerProven.

Table 51. Operating system support for ThinkSystem NVIDIA Quadro RTX 4000 8GB PCIe Active GPU, 4X67A14934

Operating systems	SR630 (Gen 2)	SR650 (Gen 2)	ST550 (Gen 2)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y
Red Hat Enterprise Linux 7.5	N	N	N	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	N	N	Y	N	N
Red Hat Enterprise Linux 8.0	N	N	N	Y	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	N	N	Y	N

ThinkSystem NVIDIA Quadro P4000 GPU

The NVIDIA Quadro P4000 is powered by NVIDIA Pascal, brings a whole new level of performance and innovative capabilities to visual computing on the desktop. Create revolutionary products. Design ground-breaking architecture. Run the most complex simulations. Tell spectacularly vivid stories in VR. The NVIDIA Quadro P4000 is certified with a broad range of sophisticated professional applications, tested by leading workstation manufacturers, and backed by a global team of support specialists.

Introduction

The NVIDIA Quadro P4000 combines a 1792 CUDA core Pascal GPU, large 8 GB GDDR5 memory and advanced display technologies to deliver the performance and features that are required by demanding professional applications. The ability to create an expansive visual workspace of up to four 5K displays (5120x2880 @ 60Hz) with HDR color support lets you view your creations in stunning detail. The P4000 is specially designed with the performance that is necessary to drive immersive VR environments. Additionally, you can create massive digital signage solutions of up to thirty-two 4K displays per system by connecting multiple P4000s via Quadro Sync II2.

Figure 13. ThinkSystem NVIDIA Quadro P4000 GPU

Part numbers

Table 52. Ordering information

Part number	Feature code	Description
4V17A10255	B225	ThinkSystem NVIDIA Quadro P4000

Technical specifications

The following table lists the specifications of the NVIDIA Quadro P4000 GPU.

Table 53. Technical specifications

Feature	Specification
GPU Memory	8 GB GDDR5
Memory Interface	256-bit
Memory Bandwidth	Up to 243 GB/s
NVIDIA CUDA Cores	1792
System Interface	PCI Express 3.0 x16
Max Power Consumption	105 W
Thermal Solution	Active
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5, DirectX 12.0, Vulkan 1.0
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 54. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro P4000, 4V17A10255	N	N	N	N	N	2	N	N	N	N	1	3	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 55. Operating system support for ThinkSystem NVIDIA Quadro P4000 8GB PCIe Active GPU, 4V17A10255

Operating systems	SR630 (Gen 2)	SR650 (Gen 2)	ST550 (Gen 2)	SR630 (Gen 1)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	Y	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 6.10	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 6.9	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.3	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.5	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	N	N	Y	N	N	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP3	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	N	N	N	Y	N	N

ThinkSystem NVIDIA Quadro P2000 GPU

ThinkSystem Quadro P2000 is the perfect balance of performance, compelling features, and compact form factor delivering incredible creative experience and productivity across a variety of professional 3D applications.

Introduction

ThinkSystem Quadro P2000 features a Pascal GPU with 1024 CUDA cores, large 5 GB GDDR5 on-board memory, and the power to drive up to four 5K (5120x2880 @ 60Hz) displays natively. This makes it an excellent choice for accelerating product development and content creation workflows that demand fluid interactivity with large, complex 3D models and scenes.

Figure 14. ThinkSystem NVIDIA Quadro P2000 GPU

Part numbers

Table 56. Ordering information

Part number	Feature code	Description
7C57A02877	B15N, AUW6*	ThinkSystem NVIDIA Quadro P2000

* Feature code varies by server

Technical specifications

The following table lists the specifications of the NVIDIA Quadro P2000 GPU.

Table 57. P2000 specifications

Feature	Specification
GPU Memory	5 GB GDDR5X
Memory Interface	160-bit
Memory Bandwidth	Up to 140 GB/s
NVIDIA CUDA® Cores	1024
System Interface	PCI Express 3.0 x16
Max Power Consumption	75 W
Thermal Solution	Active
Form Factor	4.4”H x 7.9” L, Single Slot
Display Connectors	4x DP 1.4
Max Simultaneous Displays	4 direct, 4 DP1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 120Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.53, DirectX 12.04, Vulkan 1.03
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 58. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro P2000, 7C57A02877	N	N	1	N	N	2	N	N	N	N	1	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 59. Operating system support for ThinkSystem NVIDIA Quadro P2000 PCIe Active GPU, 7C57A02877

Operating systems	ST250	SR630 (Gen 2)	ST550 (Gen 2)	SR630 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	N	Y	Y	Y	Y
Red Hat Enterprise Linux 6.10	N	N	N	Y	Y
Red Hat Enterprise Linux 6.9	N	N	N	Y	Y
Red Hat Enterprise Linux 7.3	N	N	N	Y	Y
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y
Red Hat Enterprise Linux 7.5	Y	N	N	Y	Y
Red Hat Enterprise Linux 7.6	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	N	N	Y	N	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	Y	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP4	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	N	N	N	Y	N

ThinkSystem NVIDIA Quadro P620 GPU

The NVIDIA Quadro P620 is certified with a broad range of sophisticated professional applications, tested by leading workstation manufacturers, and backed by a global team of support specialists.

Introduction

The NVIDIA Quadro P620 workstation graphics board is powered by NVIDIA Pascal GPU technology and 2 GB of GDDR5 memory. The Quadro P620 graphics board is targeted for professional CAD, DCC and visualization designers, engineers and users. Get the budget friendly highest performing graphics board to drive today's demanding professional workflows in a compact footprint. Flexible single-slot and low-profile form factor makes this card compatible with even the most space and power constrained workstation chassis.

Figure 15. ThinkSystem NVIDIA Quadro P620 GPU

Part numbers

Table 60. Ordering information

Part number	Feature code	Description
4X67A11584	B31D	ThinkSystem NVIDIA Quadro P620

Technical specifications

The following table lists the specifications of the NVIDIA Quadro P620 GPU.

Table 61. Technical specifications

Feature	Specifications
GPU Memory	2 GB GDDR5
Memory Interface	128-bit
Memory Bandwidth	Up to 80 GB/s
NVIDIA CUDA Cores	512
System Interface	PCI Express 3.0 x16
Max Power Consumption	40 W
Thermal Solution	Active
Display Connectors	4x mDP 1.4
Max Simultaneous Displays	4 direct, 4 DP 1.4 Multi-Stream
Display Resolution	4x 4096x2160 @ 60Hz 4x 5120x2880 @ 60Hz
Graphics APIs	Shader Model 5.1, OpenGL 4.5, DirectX 12.0, Vulkan 1.0
Compute APIs	CUDA, DirectCompute, OpenCL

Server support

The following table lists the ThinkSystem servers that are compatible. The numbers listed in the server columns represent the number of GPUs supported.

Table 62. ThinkSystem server support

	E	1S Intel				2S Intel								AMD		4S Intel			Dense/ Blade			
Description & Part number	SE350 (7Z46 / 7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	ST550 (7X09/7X10)	SR530 (7X07/7X08)	SR550 (7X03/7X04)	SR570 (7Y02/7Y03)	SR590 (7X98/7X99)	SR630 (7X01/7X02)	SR650 (7X05/7X06)	SR670 (7Y36/7Y37/7Y38)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR850 (7X18/7X19)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
NVIDIA Quadro P620, 4X67A11584	N	1	1	N	1	2	N	N	N	N	3	3	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 63. Operating system support for ThinkSystem NVIDIA Quadro P620 PCIe Active GPU, 4X67A11584

Operating systems	SR250	ST250	ST50	SR630 (Gen 2)	SR650 (Gen 2)	ST550 (Gen 2)	SR630 (Gen 1)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	N	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.5	Y	Y	Y	N	N	N	Y	Y	Y
Red Hat Enterprise Linux 7.6	N	N	N	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP3	Y	Y	N	N	N	N	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4	N	N	N	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	N	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	N	N	N	N	N	Y	N	N

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.
1009 Think Place - Building One
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.

This document, LP0768, was created or updated on August 20, 2019.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  http://lenovopress.com/LP0768
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at http://lenovopress.com/LP0768.

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:

Intel® and Xeon® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux® is a trademark of Linus Torvalds in the United States, other countries, or both.

DirectX®, Microsoft®, 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.