

Overview

NVIDIA Accelerators for HPE ProLiant Servers

Hewlett Packard Enterprise supports, on select HPE ProLiant servers, computational accelerator modules based on NVIDIA® Tesla™, NVIDIA® GRID™, and NVIDIA® Quadro™ Graphical Processing Unit (GPU) technology.

The NVIDIA accelerators enable seamless integration of GPU computing with HPE ProLiant servers for high-performance computing, large data center graphics and virtual desktop deployments. HPE ProLiant servers seamlessly integrate GPU computing with select HPE server families. Designed for power-efficient, high-performance supercomputing, NVIDIA accelerators deliver dramatically higher application acceleration than a CPU-only approach for a range of deep learning, scientific, and commercial applications. The thousands of NVIDIA CUDA® cores of each accelerator allow it to divide large computing or graphics tasks into thousands of smaller tasks that can be run concurrently, thus enabling much faster simulations and improved graphics fidelity for extremely demanding 3D models.

- Increased Performance to Solve Problems Faster
 - The NVIDIA accelerators for HPE ProLiant servers improve computational performance, dramatically reducing the completion time for parallel tasks, offering quicker time to solutions.
 - Co-locating the NVIDIA Quadro® or NVIDIA GRID GPUs with computational servers, large data sets can be shared, dramatically improving display refresh rates.
 - These GPUs are specifically designed to enable rich graphics in virtualized environments. Hewlett Packard Enterprise can satisfy NVIDIA GRID software via HPE Complete.
 - NVIDIA accelerators can be configured and monitored by HPE Insight Cluster Management Utility (CMU). HPE Insight CMU monitors and displays GPU health and temperature, as well as installs and provisions the GPU drivers and CUDA software.
-

Overview

NVIDIA Accelerators Models

HPE NVIDIA Quadro P2200 Graphics Accelerator	R2U55A
HPE NVIDIA Quadro P2200 Graphics Accelerator	R2U55C
HPE NVIDIA Quadro P1000 Graphics Accelerator	R3K70A
HPE NVIDIA Quadro P1000 Graphics Accelerator	R3K70C
HPE NVIDIA Quadro RTX 4000 Graphics Accelerator	R1F95A
HPE NVIDIA Quadro RTX 4000 Graphics Accelerator	R1F95C
HPE NVIDIA Quadro RTX 8000 Graphics Accelerator	R1F97A
HPE NVIDIA Quadro RTX 8000 Graphics Accelerator	R1F97C
HPE NVIDIA Quadro RTX x16 2-way 2-slot NVLink Bridge	R1F96A
HPE NVIDIA Quadro RTX x16 2-way 2-slot NVLink Bridge	R1F96C
HPE NVIDIA Quadro RTX 6000 Graphics Accelerator	R0Z45A
HPE NVIDIA Quadro RTX 6000 Graphics Accelerator	R0Z45C
HPE NVIDIA Tesla T4 16GB Computational Accelerator	R0W29A
HPE NVIDIA Tesla T4 16GB Computational Accelerator	R0W29C
HPE NVIDIA Tesla V100 PCIe 32GB Computational Accelerator	Q9U36A
HPE NVIDIA Tesla V100 SXM2 32GB Computational Accelerator	Q9U37A
HPE NVIDIA Quadro P2000 Graphics Accelerator	Q0V77A
HPE NVIDIA Quadro P4000 Graphics Accelerator	Q0V78A
HPE NVIDIA Tesla P40 24GB Computational Accelerator	Q0V80A
HPE NVIDIA Tesla P40 24GB Computational Accelerator	Q0V80C
NVIDIA Tesla M10 Quad GPU Module	Q0J62A
NVIDIA Tesla M10 Quad GPU Module	Q0J62C

NOTE: Please see the HPE ProLiant server QuickSpecs for the supported servers for configuration rules, including requirements for enablement kits.

NOTE: * CPU Type

1st Gen = 1st Gen Intel Xeon Scalable Processor

2nd Gen = 2nd Gen Intel Xeon Scalable Processor

Standard Features

Description	HPE NVIDIA Quadro P2200 Graphics Accelerator
SKU	R2U55A R2U55C
Performance	3.822 TF SP
Memory Size	5 GB GDDR5x
Memory Bandwidth	200 GB/s
Cores	1,280
System Interface	PCIe 3.0 x 16
Power	120W
Form Factor	4.4" H x 9.5" L, Single Slot

Description	HPE NVIDIA Quadro P1000 Graphics Accelerator
SKU	R3K70A R3K70C
Performance	7.1TF SP
Memory Size	4 GB GDDR5
Memory Bandwidth	Up to 80 GB/s
Cores	640
System Interface	PCIe 3.0 x 16
Power	47W
Form Factor	4.376"x 6.6" L, Single Slot

Standard Features

Description	HPE NVIDIA Quadro RTX 4000 Graphics Accelerator
SKU	R1F95A R1F95C
Performance	7.1TF SP
Memory Size	8 GB GDDR6
Memory Bandwidth	Up to 416 GB/s
Cores	2,304
System Interface	PCIe 3.0 x 16
Power	160W
Form Factor	4.4" H x 9.5" L, Single Slot

Supported Servers and Operating Systems

R1F96C Supported Servers	CPU *	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	2 nd Gen	7.6	15	NA	Win 16	NA	NA
HPE ProLiant ML350 Gen10	2 nd Gen	7.6	15	16.04	Win 16	NA	NA
HPE Synergy 480 Gen10	2 nd Gen	NA	NA	NA	NA	NA	NA
HPE ProLiant ML110 Gen10	2 nd Gen	NA	NA	NA	NA	NA	NA

Description	HPE NVIDIA Quadro RTX 8000 Graphics Accelerator
SKU	R1F97A R1F97C
Performance	16.3TF SP
Memory Size	48 GB GDDR6
Memory Bandwidth	Up to 672 GB/s
Cores	4,608
System Interface	PCIe 3.0 x 16
Power	295W
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height

Supported Servers and Operating Systems

R1F97A Supported Servers	CPU *	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant xl270d Gen10	2 nd Gen	7.5,7.6	15	18.04	Win 16	6.7	7.1
HPE Superdome Flex	2 nd Gen	NA	NA	NA	NA	NA	NA
R1F97C Supported Servers	CPU *	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	2 nd Gen	7.6	15	18.04	Win 16	6.7	7.1
HPE ProLiant DL385 Gen10	EPYC	7.6	15	NA	Win 16	6.7	7.1
HPE ProLiant xl270d Gen10	CLX	7.5, 7.6	15	18.04	Win 16	6.7	7.1

Standard Features

Description	HPE NVIDIA Quadro RTX x16 2-way 2-slot NVLink Bridge						
SKU	R1F96A R1F96C						
Supported Servers and Operating Systems							
Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant ML350 Gen10	2 nd Gen	See HPE NVIDIA RTX 6000 Graphics Accelerator for details					
HPE ProLiant ML350 Gen10	2 nd Gen	See HPE NVIDIA RTX 8000 Graphics Accelerator for details					

Description	HPE NVIDIA Quadro RTX 6000 Graphics Accelerator						
SKU	R0Z45A R0Z45C						
Performance	16.3TF SP						
Memory Size	24 GB GDDR6						
Memory Bandwidth	Up to 672 GB/s						
Cores	4,608						
System Interface	PCIe 3.0 x 16						
Power	295W						
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height						

Supported Servers and Operating Systems							
R0Z45A Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE Superdome Flex	2 nd Gen	NA	NA	NA	NA	NA	NA
HPE ProLiant xL270d Gen10	2 nd Gen	7.5, 7.6	15	18.04	Win 16	6.7	7.1
R0Z45C Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	1 st Gen	7.5	12 SP3	NA	Win 16	NA	NA
HPE ProLiant DL380 Gen10	2 nd Gen	7.5	15	16.04	Win 16	6.7	7.1
HPE ProLiant ML350 Gen10	2 nd Gen	7.5	15	NA	Win 16	NA	NA
HPE ProLiant DL385 Gen10	EPYC	7.6	15	NA	Win 16	6.7	7.1

Description	HPE NVIDIA Tesla T4 16GB Computational Accelerator						
SKU	R0W29A R0W29C						
Performance	8.1TF SP						
Memory Size	16 GB GDDR6						
Memory Bandwidth	Up to 320 GB/s						
Cores	2,560						
System Interface	PCIe 3.0 x 16						
Power	70W						
Form Factor	4.4" H x 10.5" L, Single Slot, Low Profile						

Standard Features

Supported Servers and Operating Systems							
R0W29A Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant xl190r Gen10	2 nd Gen	7.5, 7.6	15	NA	Win 16	6.7	7.1
HPE ProLiant xl270d Gen10	2 nd Gen	7.5, 7.6	15	18.04	Win 16	6.7	7.1
HPE Edgeline EL1000 m510	Xeon D	NA	15	NA	Win 10	6.7	7.1
HPE Edgeline EL1000 m710	1 st Gen	NA	15	NA	Win 10	6.7	7.1
HPE Edgeline EL4000 m510	Xeon D	NA	15	NA	Win 10	6.7	7.1
HPE Edgeline EL4000 m710	1 st Gen	NA	15	NA	Win 10	6.7	7.1
HPE Edgeline EL8000 1U e910	2 nd Gen	7.6	NA	NA	NA	6.7	7.1
HPE Edgeline EL8000 2U e910	2 nd Gen	7.6	NA	NA	NA	6.7	7.1
R0W29C Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	1 st Gen	7.5	12 SP3	16.04	Win 16	6.7	7.1
HPE ProLiant DL380 Gen10	2 nd Gen	7.5	15	16.04	Win 16	6.7	7.1
HPE ProLiant DL360 Gen10	1 st Gen	7.5	12 SP3	NA	Win 16	6.7	7.1
HPE ProLiant DL360 Gen10	2 nd Gen	7.5	15	NA	Win 16	6.7	7.1
HPE ProLiant ML350 Gen10	2 nd Gen	7.5	15	NA	Win 16	6.7	NA
HPE ProLiant DL385 Gen10	EPYC	7.5	12 SP3	16.04	Win 16	6.7	NA
HPE ProLiant DL325 Gen10	EPYC	7.5	15	NA	Win 16	6.7	7.1
HPE Synergy 480 Gen10	2 nd Gen	NA	NA	NA	NA	NA	NA

Description	HPE NVIDIA Tesla V100 SXM2 32GB Computational Accelerator
SKU	Q9U37A
Performance	7.8TF DP, 15.7TF SP, 125TF FP16
Memory Size	32GB HBM2
Memory Bandwidth	900 GB/s
Cores	5,120
System Interface	NVLink
Power	300W
Form Factor	SXM2

Supported Servers and Operating Systems							
Q9U37A Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant xl270d Gen10	1 st Gen	7.4	12 SP3	16.04	NA	NA	NA
HPE ProLiant xa780i Gen10	1 st Gen	7.4	12 SP3	NA	NA	NA	NA
HPE ProLiant xl270d Gen10	2 nd Gen	7.5	15	16.04	Win 16	6.7	7.1

Standard Features

Description	HPE NVIDIA Tesla V100 PCIe 32GB Computational Accelerator
SKU	Q9U36A Q9U36C
Performance	7TF DP, 14TF SP, 112TF FP16
Memory Size	32GB HBM2
Memory Bandwidth	900 GB/s
Cores	5,120
System Interface	PCIe Gen3
Power	250W
Form Factor	10.5"x1.37"x4.37", Full Height, Full Length

Supported Servers and Operating Systems

Q9U36A Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant xL270d Gen10	1 st Gen	7.4	12 SP3	16.04	NA	NA	NA
HPE ProLiant xL270d Gen10	2 nd Gen	7.5	15	16.04	Win 16	6.7	7.1
HPE ProLiant XL190r Gen10	1 st Gen	7.4	12 SP3	NA	Win 16	NA	NA
HPE ProLiant XL190r Gen10	2 nd Gen	7.5	15	NA	Win 16	6.7	7.1
HPE Superdome Flex	1 st Gen	7.5	15	NA	NA	NA	NA
HPE Superdome Flex	2 nd Gen	7.6	15	NA	Win 16	6.7	7.1
HPE Edgeline 8000 2U e910	2 nd Gen	7.6	NA	NA	NA	6.7	7.1
Q9U36C Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	1 st Gen	7.4	12 SP3	NA	Win 16	NA	NA
HPE ProLiant DL380 Gen10	2 nd Gen	7.5	15	16.04	Win 16	6.7	7.1
HPE ProLiant DL385 Gen10	EPYC	7.4	12 SP3	NA	Win 16	NA	NA
HPE ProLiant DL580 Gen10	1 st Gen	7.4	12 SP3	NA	Win 16	6.5 U2	NA
HPE ProLiant DL580 Gen10	2 nd Gen	7.5	15	NA	Win 16	6.7	7.1
HPE Cloudline CL2800 Gen10	1 st Gen	7.5	12 SP3	NA	Win 16	6.5 U2	7.5

Description	HPE NVIDIA Tesla M10 Quad GPU Module
SKU	Q0J62A Q0J62C
Memory Size	32GB GDDR5
Cores	2,560
System Interface	PCIe Gen3
Power	225W
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height

Standard Features

Supported Servers and Operating Systems							
Q0J62A Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant xl190r Gen10	1 st Gen	NA	NA	NA	NA	NA	NA
HPE ProLiant xl190r Gen10	2 nd Gen	NA	NA	NA	NA	6.7	7.1
Q0J62C Supported Servers	CPU	RHEL Enterprise	SLES Enterprise	Ubuntu Community	Windows Enterprise	VMware ESXi	Citrix XenServer
HPE ProLiant DL380 Gen10	1 st Gen	NA	NA	NA	NA	NA	NA
HPE ProLiant DL380 Gen10	2 nd Gen	NA	NA	NA	NA	6.7	7.1
HPE ProLiant DL385 Gen10	EPYC	NA	NA	NA	NA	6.5 U2	7.5
HPE ProLiant DL350 Gen10	1 st Gen	NA	NA	NA	NA	NA	NA
HPE ProLiant DL350 Gen10	2 nd Gen	NA	NA	NA	NA	6.7	7.1

Performance of the Quadro, P2000, P4000, and P6000 24GB Graphic Accelerators

- P2000 has 1024 cores, P4000 has 1792 cores, and P6000 has 3840 cores.
- P6000 24GB has 24 GB GDDR5 memory
- Power consumption: 75W, 105W, 250W
- Support OpenGL 4.3, Shader Model 5.0, DirectX 11
- Dedicated H.264 encode engine that's independent of 3D/compute pipeline and delivers real-time performance for transcoding, video editing, and other encoding applications.
- Provides the ability to texture from and render to 16K x 16K surfaces. This is beneficial for applications that demand the highest resolution and quality image processing.
- NVIDIA SMX delivers more processing performance and efficiency through a new, innovative streaming multiprocessor design that allows a greater percentage of space to be applied to processing cores versus control logic, enabling greater model complexity.
- Hyper-Q feature that enables multiple CPU cores to simultaneously utilize the CUDA cores on a single GPU.
- The high speed PCIe Gen 3.0 data transfer, available on the, P4000 (8GB), and P6000 (24GB) maximizes bandwidth between the HPE ProLiant server and the Tesla processors.

Performance of the Tesla P40 24GB Computational Accelerator

- The Tesla P40 is for single-precision, especially deep-learning, applications
- 3840 CUDA cores
- GPU Boost enables opportunistic clock frequency bursts provided no thermal or power limits are hit
- 12 Tflops of single-precision peak performance
- Power consumption: 250W
- 24 GB of GDDR5 memory optimizes performance and reduces data transfers by keeping large data sets in local memory
- The NVIDIA Parallel DataCache™ accelerates algorithms where data addresses are not known beforehand. This includes a configurable L1 cache per Streaming Multiprocessor block and a unified L2 cache for all of the processor cores.
- Asynchronous transfer turbo charges system performance by transferring data over the PCIe bus while the computing cores are crunching other data. Even applications with heavy data-transfer requirements, such as seismic processing, can maximize the computing efficiency by transferring data to local memory before it is needed.
- Dynamic Parallelism capability that enables GPU threads to automatically spawn new threads.
- Hyper-Q feature that enables multiple CPU cores to simultaneously utilize the CUDA cores on a single GPU.
- The high speed PCIe Gen 3.0 data transfer maximizes bandwidth between the HPE ProLiant server and the Tesla processors.

Service and Support

Service and Support

If this is a qualified option, it is covered under the HPE Support Service(s) applied to the HPE ProLiant Server. Please check HPE ProLiant Server documentation for more details on the services for this particular option.

Warranty and Support Services

Warranty and Support Services will extend to include HPE options configured with your server or storage device. The price of support service is not impacted by configuration details. HPE sourced options that are compatible with your product will be covered under your server support at the same level of coverage allowing you to upgrade freely. Installation for HPE options is available as needed. To keep support costs low for everyone, some high value options will require additional support. Additional support is only required on select high value workload accelerators, fibre switches, InfiniBand and UPS batteries over 12KVA. See the specific high value options that require additional support [HERE](#)

Protect your business beyond warranty with HPE Support Services

HPE Pointnext provides a comprehensive portfolio including Advisory and Transformational, Professional, and Operational Services to help accelerate your digital transformation. From the onset of your transformation journey, Advisory and Transformational Services focus on designing the transformation and creating a solution roadmap. Professional Services specializes in creative configurations with flawless and on-time implementation, and on-budget execution. Finally, operational services provides innovative new approaches like Flexible Capacity and Datacenter Care, to keep your business at peak performance. HPE is ready to bring together all the pieces of the puzzle for you, with an eye on the future, and make the complex simple.

Parts and materials

Hewlett Packard Enterprise will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

The defective media retention service feature option applies only to Disk or eligible SSD/Flash Drives replaced by Hewlett Packard Enterprise due to malfunction.

HPE Support Center

The HPE Support Center is a personalized online support portal with access to information, tools and experts to support HPE business products. Submit support cases online, chat with HPE experts, access support resources or collaborate with peers.

Learn more: <http://www.hpe.com/support/hpesc>

HPE's Support Center Mobile App* allows you to resolve issues yourself or quickly connect to an agent for live support. Now, you can get access to personalized IT support anywhere, anytime.

HPE Insight Remote Support and HPE Support Center are available at no additional cost with a HPE warranty, HPE Support Service or HPE contractual support agreement.

NOTE: *HPE Support Center Mobile App is subject to local availability.

For more information

Visit the Hewlett Packard Enterprise Service and Support [website](#).

Summary of Changes

Date	Version History	Action	Description of Change
02-Dec-2019	Version 31	Changed	Overview and Standard Features sections were updated Obsolete SKUs were removed
03-Jun-2019	Version 30	Changed	Models and Standard Features sections were updated
02-Apr-2019	Version 29	Changed	Overview and Standard Features sections were updated. Obsolete SKUs were removed.
04-Feb-2019	Version 28	Changed	Overview and Standard Features sections were updated
03-Dec-2018	Version 27	Changed	Standard Features section was updated
15-Oct-2018	Version 26	Changed	Overview and Standard Features sections were updated.
01-Oct-2018	Version 25	Changed	Overview, Additional Options and Service and Support sections were updated. Added GV100 Bridge Kit; corrected wrong PNs
04-Jun-2018	Version 24	Changed	Changes made throughout document; adding V100 FHHL card
07-May-2018	Version 23	Changed	Add HPE NVIDIA Tesla V100 PCIE 32GB Computational Accelerator / HPE NVIDIA Tesla V100 SXM2 32GB Computational Accelerator
02-Apr-2018	Version 22	Changed	Updates throughout document; added GV100
04-Dec-2017	Version 21	Changed	New modules added
16-Oct-2017	Version 20	Changed	Update verbiage in the Standard
25-Sep-2017	Version 19	Changed	Update verbiage in the Standard Features Sections
11-Jul-2017	Version 18	Changed	New models were added and updates throughout the whole document
05-Jun-2017	Version 17	Changed	Remove obsolete info and SKUs and update the recent info
08-May-2017	Version 16	Changed	New models were added and updates throughout the whole document
28-Nov-2016	Version 15	Changed	Updates throughout the whole document
26-Sep-2016	Version 14	Changed	Added new SKUs to QS
6-Jun -2016	Version 13	Changed	Add The new information to the QS and remove obsolete SKU 's
31-Mar-2016	Version 12	Changed	Update all the sections and general info throughout the QuickSpecs
01-Dec-2015	Version 11	Changed	Update the Standard Features and the technical Specifications section
17-Aug-2015	Version	Changed	Update several Overview and technical specifications.
09-Feb-2015	Version 9	Changed	Update several Overview and technical specifications.
01-Dec-2014	Version 8	Changed	Revised wording and Technical Specifications
09-Sep-2014	Version 7	Changed	Changes made throughout the QuickSpecs.
05-Jun-2014	Version 6	Changed	High Performance Clusters and Thermal Solutions were revised
31-Mar-2014	Version 5	Changed	NVIDIA Tesla K40C 12 GB Computational Accelerator and NVIDIA Quadro K2000 PCIe Graphics Adapter were added
18-Feb-2014	Version 4	Changed	Changes made throughout the QuickSpecs
09-Dec-2013	Version 3	Changed	NVIDIA Tesla K10 Rev B Dual GPU Module and NVIDIA Tesla K40 12 GB Module were added.
20-Sep-2013	Version 2	Changed	Changes made in the following Sections Standard Features Optional Features Technical Specifications
10-Sep-2013	Version 1	New	New QuickSpecs

Summary of Changes



© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Windows and Microsoft are registered trademarks of Microsoft Corp, in the U.S.

c04123180 - 14576 - WorldWide - V31 - 02-December-2019