

Family data sheet



Accelerators for HP ProLiant servers

Enable scalable and efficient high-performance computing



November 2014

HP high-performance computing has made it possible to accelerate innovation at any scale. But traditional CPU technology is no longer capable of sufficiently scaling performance to address the skyrocketing demand for compute resources. HP high-performance computing solutions are built on HP ProLiant servers using industry-leading accelerators to dramatically increase performance with lower power requirements.

Innovation is the foundation for success

What is hybrid computing?

A hybrid computing model is one where accelerators (known as GPUs or coprocessors) work together with CPUs to perform computing tasks.

As parallel processors, accelerators can split computations into hundreds or thousands of pieces and calculate them simultaneously.

Offloading the most compute-intensive portions of applications to accelerators dramatically increases both application performance and computational efficiency.

Accelerators are revolutionizing high performance computing

High-performance computing (HPC) is being used to address many of modern society's biggest challenges, such as designing new vaccines and genetically engineering drugs to fight diseases, finding and extracting precious oil and gas resources, improving financial instruments, and designing more fuel efficient engines.

This rapid pace of innovation has created an insatiable demand for compute power. At the same time, multiple strict requirements are placed on system performance, power consumption, size, response, reliability, portability, and design time. Modern HPC systems are rapidly evolving, already reaching petaflop and targeting exaflop performance.

All of these challenges lead to a common set of requirements: a need for more computing power with increased power efficiency at higher densities.

Accelerated computing builds on the power of HPC, by using accelerators—called graphics processing units (GPU) or coprocessors—for parallel processing in tandem with CPUs in a hybrid computing model. By offloading the most compute-intensive portions of the application to the accelerators and running the remainder on the CPU, hybrid computing allows you to achieve unprecedented application performance. In addition, accelerators deliver greater performance per watt of power consumed for a superior performance/power ratio.

HP leads the way for accelerated computing

HP has been a driving force in bringing accelerated computing into HPC, just as we led the transition to industry-standard clusters in the last decade. In fact, HP was the first company to build an industry-standard server with integrated NVIDIA general purpose GPUs. We are now shipping our 5th generation of ProLiant servers with integrated accelerators.

Today, HP is a leading provider of accelerated computing solutions with NVIDIA® GPUs and Intel® coprocessors. Our solutions combine the power of an industry-leading, high-performance computing infrastructure with highly efficient systems and solutions to help you solve scientific, engineering, and data analysis problems—empowering innovation at any scale, with the affordability of volume, industry-standard technologies.

HP accelerated computing solutions allow you to work:

Faster—Speed advancements with converged infrastructure that is purpose-built for scale.

Better—Optimize your performance footprint with the world's most efficient systems.

Smarter—Deploy easily, adapt quickly to change, and improve quality of service.

“With this supercomputer based on HP ProLiant...we have doubled the performance using the same amount of power and the same amount of floor space as our previous clusters.”

—Mike Shuey, HPC Systems Manager, Purdue University

HP ProLiant servers for accelerated computing

HP offers a broad spectrum of innovative, industry standard HPC solutions based on HP ProLiant servers with support for industry-leading accelerators from NVIDIA and Intel. Our solution offerings include:

- Fast and simple clustering—HP Cluster Platform combines the flexibility of a custom solution with the simplicity, reliability, and value of a preconfigured, factory-built product
- A breadth of servers leveraging accelerators—The HP accelerator-enabled server portfolio includes HP Apollo 6000 and HP Apollo 8000 System servers, HP ProLiant SL servers, and HP ProLiant DL and ML servers
- Accelerator-aware cluster management—HP Insight Cluster Management Utility (CMU)
- Experience and expertise—Over a decade of development and support, with customer deployment at global sites

Our broad portfolio delivers superior server choice that reliably supports any workload. Our focus on innovation allows you to do more with less to reduce your capital and operating expenses and maximize return-on-investment. Our deep experience and expertise let you select and deploy your solution with confidence.

Server portfolio

- HP Apollo 8000 System servers

A supercomputer system combining high-level processing power with groundbreaking warm-water cooling for ultra-low energy usage and recycling

- HP Apollo 6000 System servers

High-density, air-cooled system optimizes rack-scale performance, and makes cost-effective HPC more accessible

- HP ProLiant SL family servers

Shared infrastructure servers ideal for maximum HPC scale, efficiency, and density

- HP ProLiant DL family servers

Versatile rack-optimized servers ideal for flexible HPC configurations

- HP ProLiant ML family servers

Expandable tower servers that are ideal for remote and branch offices

- HP ProLiant WS460c Gen8 graphics server blade

The industry's most comprehensive graphics acceleration capabilities for VDI in a blade form factor. See the [data sheet](#) for more information.

Case in point: The TSUBAME 2.5 supercomputer

The Tokyo Institute of Technology's TSUBAME 2.5 supercomputer was created by upgrading the TSUBAME 2.0 system's 1,408 HP ProLiant SL servers from older GPUs to the latest NVIDIA Tesla GPUs. The upgrade increased peak performance from 2.4 petaflops to 5.7 petaflops and LINPACK performance from 1.19 petaflops to 2.78 petaflops. These performance improvements were achieved without increasing the system's data center footprint, while reducing the operational power consumption of the machine by 18 percent.¹

¹ HP case study, “A Green Machine – Tokyo Institute of Technology proves that supercomputers can be both fast and energy efficient,” May 2014

Integrated accelerator solutions for HP ProLiant Servers

NVIDIA GPUs

Based on NVIDIA's CUDA architecture, NVIDIA accelerators enable seamless integration of GPU computing with HP ProLiant servers for high-performance computing, large data center graphics, and virtual desktop deployments. These accelerators deliver all of the standard benefits of GPU computing while enabling maximum reliability and tight integration with system monitoring and management tools like HP Insight CMU.



- NVIDIA® Tesla® GPUs excel at boosting performance of structured numerical algorithms along with power efficiency. Applications that benefit from these accelerators include seismic processing, biochemistry simulations, weather and climate modeling, image, video and signal processing, computational finance, computational physics, CAE, CFD, and data analytics.

– For single-precision algorithms

NVIDIA Tesla K10 Dual GPU 8 GB and NVIDIA Tesla K80 Dual GPU 24 GB

– For double-precision algorithms

NVIDIA Tesla K20 5 GB, NVIDIA Tesla K20X 6 GB, NVIDIA Tesla K40 12 GB, NVIDIA Tesla K80 Dual GPU 24 GB, NVIDIA Tesla K40 XL dense form factor



- NVIDIA® Quadro® GPUs offer outstanding graphics performance on a range of professional applications including remote visualization with multi-monitor capability and large-scale and high-resolution 3D remote visualization.

– For remote visualization with multi-monitor capability

NVIDIA Quadro K2200 and NVIDIA Quadro K4200

– For large-scale and high-resolution 3D remote visualization

NVIDIA Quadro K5200 and NVIDIA Quadro K6000



- NVIDIA® GRID™ GPUs are optimized for virtual desktop infrastructures (VDI) and support large numbers of users with standard desktop applications or virtual desktop applications all in the same data center; they can also be used to enable professional-class visualization features.

– For standard desktop applications

NVIDIA GRID K1 Quad GPU


– For professional-class visualization

NVIDIA GRID K2 Dual GPU

Since the 1990s, HP and NVIDIA have shared a deep partnership. Our shared expertise in HPC and GPUs fostered early collaboration using GPUs for computation, which became a focus area for HP's HPC innovation. The Accelerator Team within HP's HPC organization has been working with NVIDIA for over a decade, and HP has a Technical Centre of Excellence with NVIDIA located in Grenoble, France.

The NVIDIA CUDA programming environment

CUDA is NVIDIA's parallel computing architecture. The CUDA programming environment has broad support of programming languages and application programming interfaces (APIs). Choose C, C++, OpenCL, DirectCompute, or Fortran to express application parallelism and take advantage of the innovative NVIDIA GPU architectures, which are capable of running thousands of parallel threads. The CUDA programming model lets programmers focus on parallelizing algorithms instead of the mechanics of the language. The CUDA software, as well as the GPU drivers, can be automatically installed on HP ProLiant servers by HP Insight Cluster Management Utility.



HP has collaborated with Intel to ensure that supported HP ProLiant servers fully exploit the power of Intel Xeon Phi coprocessors. Intel Xeon Phi coprocessors are built to support all your mission-critical applications in a manner that saves both time and energy, letting you use your assets where you truly need them—for innovations, discoveries, and more.

Intel® Xeon Phi™ coprocessors

Based on the Intel Many Integrated Core (MIC) architecture, these coprocessors enable highly parallel computing that is quick and efficient, with more than one teraflop of double-precision peak performance in every chip. These coprocessors deliver all of the standard benefits of accelerator computing while enabling maximum reliability and tight integration with system monitoring and management tools such as HP Insight CMU.

- Intel Xeon Phi 5110P

The Intel Xeon Phi coprocessor 5100 series is optimized for high-density computing and is well-suited for workloads that are memory-bandwidth bound, such as STREAM; memory-capacity bound, such as ray tracing; or both, such as reverse time migration.

- Intel Xeon Phi 7120P and Intel Xeon Phi 7120D

The Intel Xeon Phi coprocessor 7100 series provides the most features and the highest performance and memory capacity of the Intel Xeon Phi product family. The family supports Intel Turbo Boost Technology 1.0, which increases core frequencies during peak workloads when thermal conditions allow.

Intel Cluster Studio XE development environment

Intel Xeon Phi coprocessors use the same programming model that is used on Intel Xeon processors, allowing for common source code that can execute on either Xeon processors or Xeon Phi coprocessors.

Intel Cluster Studio XE is the primary development environment for developers leveraging the Message Passing Interface standard (MPI) to develop and deploy highly parallel programs that require cluster level computations. Parallel compute models need to scale on both multi-CPU and coprocessor environments. The key in doing this is a tool set that ensures correctness and optimally utilizes all CPU and coprocessor resources, getting your app working quickly and efficiently on platforms—both today and in the future.

With alternative technologies, you typically have to re-write your program, and some applications can't be rewritten because they're too large or complex. Intel Cluster Studio XE addresses these issues. This tool set puts together the industry leading MPI library that continues to scale. It pairs this with Intel C++ and Fortran compilers, as well as performance libraries and powerful programming models. Intel Cluster Studio XE is the package that allows you to take advantage of the hardware and to write both MPI and hybrid MPI programs.

Supported accelerators by server model

HP ProLiant rack and tower servers

HP ProLiant DL580 Gen8	HP ProLiant DL380 Gen9	HP ProLiant DL360 Gen9	HP ProLiant DL180 Gen9	HP ProLiant DL120 Gen9	HP ProLiant DL80 Gen9	HP ProLiant ML350 Gen9	HP ProLiant ML150 Gen9
NVIDIA Tesla K40C	NVIDIA Tesla K40C Intel Xeon Phi 5110					NVIDIA Tesla K40C	
NVIDIA GRID K2 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU	NVIDIA Quadro K2200 GPU
NVIDIA Quadro K6000 GPU	NVIDIA Quadro K4200 GPU NVIDIA Quadro K6000 GPU NVIDIA GRID K1 GPU NVIDIA GRID K2 GPU	NVIDIA Quadro K4200 GPU	NVIDIA Quadro K4200 GPU	NVIDIA Quadro K4200 GPU	NVIDIA Quadro K4200 GPU	NVIDIA Quadro K4200 GPU NVIDIA Quadro K5200 GPU NVIDIA Quadro K6000 GPU	

HP ProLiant hyperscale servers

HP Apollo 8000 System		HP Apollo 6000 System	HP ProLiant SL6500 Scalable System		HP ProLiant SL2500 Scalable System
HP ProLiant XL740f	HP ProLiant XL750f	HP ProLiant XL250a Gen9	HP ProLiant SL250s Gen8	HP ProLiant SL270s Gen8	HP ProLiant SL210t Gen8
Intel Xeon Phi 7120D	NVIDIA Tesla K40 XL	NVIDIA Tesla K40 12 GB NVIDIA Tesla K80 24 GB Intel Xeon Phi 5110P Intel Xeon Phi 7120P	NVIDIA Tesla K10 Dual GPU NVIDIA Tesla K20 5 GB NVIDIA Tesla K20X 6 GB NVIDIA Tesla K40 12 GB Intel Xeon Phi 5110P Intel Xeon Phi 7120P	NVIDIA Tesla K10 Dual GPU NVIDIA Tesla K20 5 GB NVIDIA Tesla K20X 6 GB NVIDIA Tesla K40 12 GB Intel Xeon Phi 5110P	
			NVIDIA GRID K2 GPU	NVIDIA GRID K2 GPU	NVIDIA Quadro K2000 GPU NVIDIA Quadro K4000 GPU NVIDIA Quadro K5000 GPU NVIDIA Quadro K6000 GPU

For more information on accelerators see: hp.com/us/en/products/server-comp-graphic-accelerators/

For more information on HP ProLiant servers see: hp.com/go/proliant

 = Computational accelerators

 = Graphics accelerators

Completing your solution

HP Cluster Platform at-a-glance

The HP Cluster Platform includes:

- Codified best practices for building a cluster
- Configurator for fast and correct ordering
- Standardized, qualified implementation
- Recommended configurations for nodes, interconnects, and storage enabling HP and our partners to confidently deploy clusters in less time
- Direction for manufacturing on where to place components to enable the best airflow and serviceability

Learn more at hp.com/go/clusterplatforms

HPC is more than just a server with an accelerator. HP has a full range of offerings that allow you to design your complete HPC solution according to your workload characteristics, budget, and expertise. Start with the HP ProLiant server chassis, and then build out your HP Cluster Platform around your applications and performance requirements, built to respect your density or power constraints and budget. Complete your solution with:

HP Insight Cluster Management Utility

Management is a critical component of any accelerator-enabled HPC environment. HP Insight Cluster Management Utility (CMU) eases the burden of managing thousands of compute nodes—including both CPUs and accelerators. This unique, accelerator-aware cluster management software is designed as a lightweight, flexible management system with an intuitive graphical interface that enables you to visualize your entire cluster. And you can install the OS and system image on one or 1,000 servers, all from scratch, in less than half an hour. Learn more at hp.com/go/icmu

HP iLO Management

HP iLO Management with Integrated Lifecycle Management provides new levels of performance and quality of service with HP Active Health and Agentless Management. Monitoring the health of HPC solutions has traditionally required running monitoring software on the systems in an HPC solution, stealing cycles from the system's primary computational tasks. With Active Health and Agentless Management, all the monitoring is performed on the iLO Management processors, allowing extensive monitoring without impacting performance.

For additional capabilities, the HP iLO Advanced license offers iLO Federation features such as Group Firmware Updates, Group Virtual Media, Group Power Control, Group Power Capping and Group License Activation. You can also activate the enhanced Integrated Remote Console, advanced security features, power management functionality and many others. For additional licensing options visit hp.com/go/iLO

Customization and deployment services along with your storage and server purchases. You can customize hardware to your exact specifications in the factory—helping speed deployment. hp.com/go/factoryexpress

Explore HP server supported computational and graphics accelerators: hp.com/us/en/products/server-comp-graphic-accelerators/

Gain the skills you need with ExpertOne training and certification from HP. With HP ProLiant training, you will accelerate your technology transition, improve operational performance, and get the best return on your HP investment. Our training is available when and where you need it, through flexible delivery options and a global training capability. hp.com/learn/proliant

HP services and support

HP has a global team of award-winning HPC services experts available to help design, deploy, manage, and support your HPC environment and processes, including consulting, integration, outsourcing, and support.

HP Datacenter Care is a flexible, comprehensive, relationship-based approach to personalized support and management of heterogeneous data centers. With a structured framework of repeatable, tested, and globally available services, your team can leverage HP's experience supporting complex environments, global support partnerships, and technical expertise. You get exactly the services you need—when and where you need them—in a single agreement.

For more information on our services and support portfolio, visit: hp.com/services.

HP Financial Services Systems

Having access to technology on terms that align to your business needs is critical, and HP Financial Services is uniquely positioned to help accelerate your move to the data center of the future with a broad portfolio of flexible investment and transition solutions. Maximize your current data center environment, and access the latest high-performance computing technology when you need it with HP Financial Services. hp.com/go/hpfinancialservices

Learn more at
hp.com/go/hpc
hp.com/go/proliant
hp.com/go/apollo
hp.com/go/icmu

Sign up for updates
hp.com/go/getupdated



Share with colleagues



Rate this document

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP 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. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and/or other countries.

4AA5-5836ENW, November 2014

