

HP ProLiant SL6500 Scalable Systems with GPUs



New levels of heterogeneous computing performance

Table of contents

Introduction.....	2
New servers, new innovation, new economies of scale.....	2
HP ProLiant SL6500 Gen8 Scalable System hardware infrastructure.....	3
Comparison of the HP ProLiant SL250s and SL270s Gen8 SE Servers	4
GPU, FlexibleLOM, InfiniBand, and Ethernet connectivity	5
HP and NVIDIA.....	7
NVIDIA Tesla GPUs with Kepler architecture.....	7
HP HPC management tools	9
HP Insight Cluster Management Utility.....	9
HP SL Advanced Power Manager	9
The iLO Management Engine.....	9
Part of the HP Unified Cluster Portfolio.....	10
HP Cluster Platform at a glance.....	10
HPC partners.....	11
Summary.....	12
Resources.....	13

Introduction

The high-performance computing (HPC) “grand challenges,” long established as goals for fundamental research in the areas of science and engineering, have moved beyond the laboratory and into mainstream industries. HPC computational tools address many of today’s grand challenges; such as life sciences efforts in designing a new vaccine to avert the next flu pandemic, genetically engineering drugs to fight Alzheimer’s, generating highly effective seismic analysis, improving financial instruments, and designing more fuel efficient engines. All of these challenges have one thing in common: the need for more computing power with increased power efficiency at higher densities. Working together, HP and NVIDIA® have developed integrated heterogeneous computing systems that deliver improvements in performance, efficiency, and density. These heterogeneous HP HPC cluster systems deliver more computing power, require less power and cooling, and consume substantially less space. These systems, including the HP ProLiant SL6500 Scalable System, offer a balanced approach to take full advantage of both CPU and optional graphics processing unit (GPU) in order to maximize every hour, watt, and dollar. Since the expectation is for the GPUs to carry out a substantial portion of the calculations, the HP ProLiant SL6500 Scalable System is designed such that interconnect, host memory, PCIe bus, and network interconnect performance characteristics are matched with the GPU performance in order to maintain a well-balanced system.

Heterogeneous supercomputing, incorporating combined general-purpose and accelerated CPUs, hold promise to be the future architecture of choice due to their broad application base and superior performance: power ratio¹. No longer confined to exclusive environment of fundamental scientific research institutions, heterogeneous HPC environments are now found in a broad range of commercial applications, and academic institutions:

- Energy/Utilities - Power flow studies are one of the most important aspects of power system planning and operation. GPUs can deliver improvements measured in orders of magnitude for these simulations.
- Seismic analysis – Software providers have found that GPUs are especially effective for accelerating the processing of multi-terabytes of raw seismic data and visualizing the results.
- Communications - Antenna array design involves repeated simulation to tune the many parameters involved. Offloading the optimization workload onto GPUs reduces that time significantly.
- Financial instruments - Options pricing, risk analysis, algorithmic trading, and Monte Carlo simulation can all benefit from GPUs.
- Weather simulation - Most climate models in use today are typically run at global resolutions of 100 to 200 km. This level of resolution limits the ability to model cloud-system behaviors, GPUs enable models to run at less than 4 km.
- Genomics and Life Sciences - Sequencing and protein docking are very compute-intensive tasks that see a large performance benefit by using GPUs.

New servers, new innovation, new economies of scale

Regardless of industry or institutional focus, you need a balanced computing system to avoid bottlenecks and help crunch more numbers in less time at lower cost. Organizations are searching for an IT solution that provides more computing capability and better power efficiency, but also increases density.

HP designed the HP ProLiant SL6500 Gen8 Scalable System family to meet the specific requirements of this growing population of customers. The efficient modular architecture of the HP ProLiant SL6500 Gen8 Scalable System offers varied compute, storage, and accelerator configurations that enable clients to optimize their infrastructure for many different HPC workloads. This removes the need to assemble hardware in untested configurations, and at the same time uses fewer components and less energy for power and cooling.

Our chief design goals for the HP ProLiant SL6500 Gen8 system servers were performance, power efficiency and density. The resulting design incorporates up to eight NVIDIA Tesla® GPUs or four NVIDIA GRID® K2 boards, two Intel® Xeon® E5-2600 CPUs, and 16 DIMM slots in each half-width server, giving you a maximum GPU density of 144 GPUs, with 32 CPUs, in a 42U rack. We ensured performance by maximizing PCIe bandwidth to each GPU. And finally, we included optional InfiniBand (IB) connectivity, providing fast and reliable GPU interconnect capabilities, with the added benefit of lower cost than other available solutions.

The HP SL6500 Scalable System uses the 4U ProLiant s6500 chassis. The chassis can accommodate two, four, or eight half-width servers giving you the ability to mix and match server nodes, and also allows for single-node serviceability - making it an extremely flexible platform. The highly efficient multi-node chassis uses shared power and cooling components, including energy efficient fans and optional redundant hot-plug power supplies to help reduce power consumption compared to traditional servers. The SL6500 Scalable System offers a choice of HP ProLiant SL230s Gen8, SL250s Gen8, or

¹ “Massive supercomputing coping with heterogeneity of modern accelerators” by Toshio Endo and Satoshi Matsuoka

SL270s Gen8 SE Servers. Networking capabilities include integrated Gigabit Ethernet, plus with optional FlexibleLOM, low-latency fourteen data rate (FDR) 56Gb/s InfiniBand, and/or 10 Gb/s Ethernet are available.

To meet HP green requirements² for the ProLiant SL6500 Gen8 Scalable System, we used 94% efficient power supplies (80Plus certified Platinum), which help ease the power consumption burden imposed by the GPUs (75% of all power used by the server). You also get more compute power from every rack because each server includes more GPUs. The ProLiant S6500 chassis shared power architecture, SL Advanced Power Manager (SL-APM), and the iLO Management Engine capabilities provide high levels of power efficiency, at the server, chassis and rack level.

Regardless of your IT system design, the flexibility of the ProLiant SL6500 Scalable System lets you use every slot for your unique requirements, changing the configuration as your needs evolve. Available in three half-width server trays, the 1U HP ProLiant SL230s Gen8 server with two CPUs, the 2U HP ProLiant SL250s Gen8 Server with two CPUs and up to three GPUs, and the HP ProLiant SL270s Gen8 SE Server with two CPUs and up to eight GPUs, gives you the flexibility to choose the one that best suits your needs and budget.

HP recognized the potential of the innovative technology NVIDIA was incorporating into its Tesla GPU accelerators, which significantly boosts performance by adding the computational power of the GPUs to the CPUs. To tailor this important benefit for our HPC customers, HP formed a collaborative partnership with NVIDIA to include these capabilities in the new HP ProLiant SL250s and SL270s Gen8 SE servers. As the first vendor to take this leap in innovation, HP has taken a leadership position in heterogeneous GPU-enabled systems and elevated NVIDIA accelerators from a specialized function to a mainstream solution.

This purpose-built architecture, designed for clusters that can include thousands of nodes, helps create competitive advantage by using every watt and square foot with maximum efficiency. HP builds upon the SL6500 infrastructure to offer a complete HPC solution, based on HP's Unified Cluster Strategy. This starts by incorporating the SL6500 server family into the HP Cluster Platform offerings³, and we provide a full complement of software and services to simplify and optimize deployment.

HP ProLiant SL6500 Gen8 Scalable System hardware infrastructure

The HP SL6500 Gen8 Scalable System is our third generation of GPU-enabled servers. This well established history of deploying HPC clusters provides the experience and knowledge to produce the HP ProLiant SL6500 Gen8 Scalable System. You also get best practices from the preceding generations of HP ProLiant SL servers. The result is a series of massively scalable systems that offer integrated application and power management, very efficient power and cooling, and much more.

The HP ProLiant SL250s Gen8 Server supports up to three NVIDIA Tesla GPUs or up to three NVIDIA GRID K2 boards, offering multiple options for interconnect, I/O expandability, and networking. This enables the SL250s server to deliver excellent compute density, power efficiency and affordability. The HP ProLiant SL270s Server supports up to eight GPUs or up to four NVIDIA GRID K2 boards for the most extreme computing tasks.

The HP ProLiant SL250s and SL270s Gen8 SE servers feature dual multi-core Intel® Xeon® processors, dedicated PCIe Gen2 and 3 with 16 lanes to each GPU, optional fourteen data rate (FDR) 56Gb/s single or dual-rail Infiniband (IB) capability and, a software ecosystem to simplify deployment. These ultra-high-density servers are the platform of choice for HPC, Web 2.0 and the most intense data sites; enabling you to solve bigger problems at an affordable price.

The HP ProLiant SL6500 Gen8 Scalable System offers you the proven HP ProLiant technical foundation in a scale out, next generation, GPU-enabled server. You can choose from 1U half-width HP ProLiant SL230s Gen8 servers with no GPU option, 2U half-width HP ProLiant SL250s Gen8 servers with up to three integrated NVIDIA Tesla or up to three NVIDIA GRID K2 boards, or 4U half-width HP ProLiant SL270s Gen8 SE servers with up to eight integrated NVIDIA Tesla GPUs or up to four NVIDIA GRID K2 boards. The SL230s 1U servers offer:

- Integrated 1GbE, with the optional FlexibleLOM for multiple networking options including low-latency 10GbE or FDR IB adaptors
- HP iLO Management Engine features support the server lifecycle from deployment through ongoing management to service alerting
- Density – Eight servers in 4U
- Front serviceability
- Rich configuration – Full-power CPUs in half-width 1U with large 256GB memory capacity

² To find out more about how HP helps you 'Go Green', go to: hp.com/hho/cache/573579-0-0-225-121.html

³ You can read more about the complete family of HP Cluster offerings at: hp.com/V2/GetPDF.aspx%2F4AA0-6746ENW.pdf

HP Performance Optimized Datacenters (PODs) are emerging as an attractive alternative to traditional brick and mortar facilities. Arriving ready to run with rack-mount servers and storage, cooling, networking, cabling and power in a highly efficient self-contained enclosure, HP PODs allow computing resources to be deployed where and when needed. You can read more about HP POD solutions at hp.com/go/pod.

Today, HP is the only vendor offering common slot hot plug Platinum power supplies (80Plus certified) with 94% efficiency. All HP Common Slot power sources are UL, CE Mark Compliant, hot-plug and support redundant configurations. HP also provides advanced power management, offering control at the GPU, chassis, and system levels.

The 2U SL250s and 4U SL270s Gen8 servers offer all the above plus integrated NVIDIA Tesla GPUs or NVIDIA GRID K2 boards for substantial additional compute capabilities in addition to hot swap drives.


Comparison of the HP ProLiant SL250s and SL270s Gen8 SE Servers

The higher density of GPUs to CPU cores means more GPUs and more computing power. Each 2U half-width server can deliver four Teraflops of Linpack⁴ performance. With two half-width servers per 2U, that means four Teraflops per 1U. For the Ultra-dense 4u half-width server that goes up substantially with each server delivering over 10 Teraflops of Linpack performance, or five Teraflops per 1U. Compared to the latest quad-core CPU, HP ProLiant SL250s Gen8 servers with NVIDIA's Tesla K20x-series GPU computing processors inside delivers 3x the double precision performance compared to the previous generation Fermi-based Tesla M2090 in the same power envelope. The HP ProLiant SL6500 Gen8 Scalable System servers incorporate the HP best practices developed from shipping accelerator-enabled systems since 2006, while other companies are still working on their first generation GPU-enabled servers. We've also designed and delivered HPC clusters since 2000. HP GPU-enabled servers are fully tested and integrated at the factory. You can work with HP to design the precise system you want, delivered fully configured and ready to run.

You get integrated system and cluster management views of CPUs and GPUs from management tools such as the HP Cluster Management utility (CMU) and HP SL Advanced Power Manager (SL-APM). You will read more about [HPC management tools](#) later in this paper. Our alliance with industry leading infrastructure solution providers allow HP to offer a system foundation of debuggers, compilers, and libraries that let you customize your HP ProLiant SL server Gen8 based cluster.

Figure 1 identifies the hardware feature design of the ProLiant SL250s server.




Figure 1.



⁴ LINPACK benchmarks are HPC benchmarks that measure floating point computing power and the ability to solve real world problems

Table 1 provides a hardware and feature comparison between the three HP ProLiant SL6500 Scalable System server options.

Table 1.

			
	SL230s	SL250s	SL270s
Form factor chassis	Fully configured s6500 chassis is 4U, full-width; 200mm extension	Fully configured s6500 chassis is 4U, full-width; 200mm extension	Fully configured s6500 chassis is 4U, full-width; 200mm extension
Processor family Intel	Intel Xeon E5-2600 v2 family	Intel Xeon E5-2600 v2 family	Intel Xeon E5-2600 v2 family
Processor number	1 or 2 processors.	1 or 2 One CPU can be utilized so long as no accelerator is implemented. 2 CPUs must be used when using an accelerator.	1 or 2 One CPU can be utilized so long as no accelerator is implemented. 2 CPUs must be used when using an accelerator.
Processor speed	1.7 through 3.5 GHz	1.7 through 3.5 GHz	1.7 through 3.5 GHz
Processor core available	4, 6, 8, 10, or 12	4, 6, 8, 10, or 12	4, 6, 8, 10, or 12
Expansion slots	One PCIe Gen3 x 16 LP slot; One PCIe Gen3 x 16 FlexLOM slot	One PCIe Gen3 x 8 LP slot; One x 16 PCIe Gen3 FlexLOM slot	Eight PCIe Gen2 x 16 accelerator slots, One PCIe Gen3 x 8 LP slot, and One PCIe Gen3 x 8 FlexLOM slot
Memory slots	16 DIMM slots maximum	16 DIMM slots maximum	16 DIMM slots maximum
Memory maximum	256 GB	256 GB	256 GB
Memory type	DDR3 RDIMM or UDIMM	DDR3 RDIMM or UDIMM	DDR3 RDIMM or UDIMM
Remote management	iLO4; IPMI 2.0 and DCMI 1.0, SLAPM	iLO4; IPMI 2.0 and DCMI 1.0 SLAPM	iLO4; IPMI 2.0 and DCMI 1.0 SLAPM
Accelerators	No support	NVIDIA K10, K20, K20x, K2 GPUs	NVIDIA K10, K20, K20x, K2 GPUs
Full configuration form factor	1U, half-width	2U, half-width	4U, half-width

While the HP ProLiant SL6500 Gen8 system servers is our flagship accelerator GPU-based server, HP offers a rich portfolio of products that also support GPUs for graphics applications. Talk to your HP representative for details and ask about HP's long and growing history of GPU support and technology leadership. The latest configuration and Quick Spec information can be found at <http://h18004.www1.hp.com/products/quickspecs/Division/Division.html#13741>

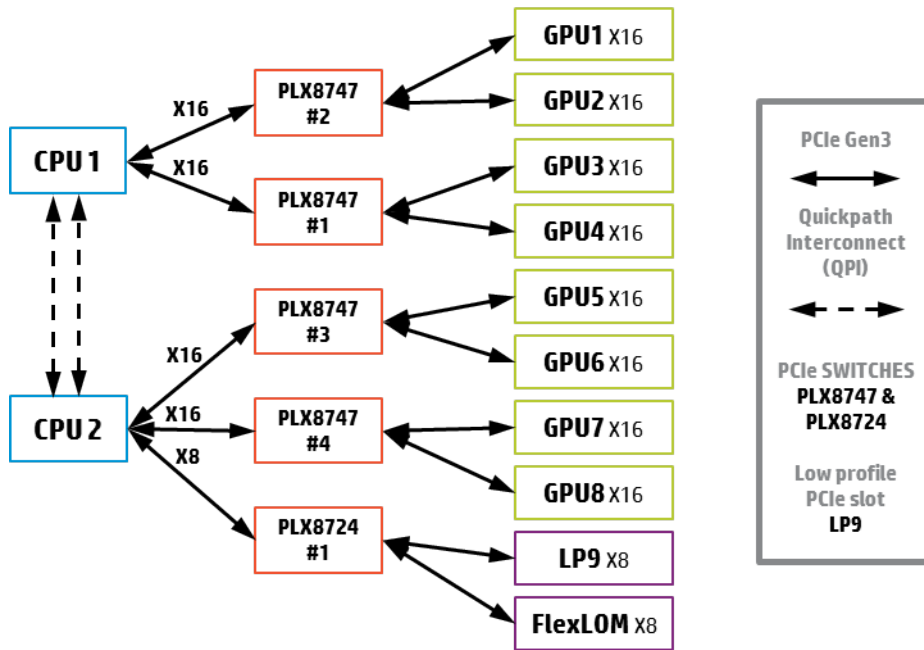
GPU, FlexibleLOM, InfiniBand, and Ethernet connectivity

HP ProLiant SL6500 Gen8 system servers have dedicated PCIe Gen2 and 3 communications lanes, with 16 lanes to each GPU. The NVIDIA K20 and K20x GPUs are Gen 2 compatible. The NVIDIA K10 and GRID K2 GPUs are PCIe Gen 3 compatible.

In addition to GPU-dedicated communication lanes, you also have more interconnect choices with the SL6500 Gen8 system servers (Figure 2). The ConnectX3 FlexibleLOM enables 40GbE SFP+ and FDR IB QSFP+ connections and the HP ProLiant SL6500 Gen8 system servers have an available x8 LP slot, and doesn't require one of the GPU slots. Many other vendors require a GPU slot be used for additional interconnects.

Figure 2.

SL270s Gen8 SE server



Servers with a NIC embedded on the system board are referred to as using LAN-on-motherboard (LOM) architecture. LOM architecture results in an efficient system design that does not require an add-on option for network functionality. However, the “one size fits all” LOM strategy can create issues. For instance, if the LOM on your server does not match your network infrastructure exactly, you will either have to settle for something less than you wanted, purchase additional hardware to make up the difference, or not use all the features that you purchased. Even if the LOM meets current needs, your network requirements could change in the future, forcing you to abandon system board components (that will continue to draw power) and add another card.

FlexibleLOM technology for HP SL Gen8 servers is a variation of LOM architecture that allows you to choose the type of NIC to be resident in the server, including ConnectX3 IB. Instead of an embedded NIC, FlexibleLOM uses a system board edge connector that accepts a FlexibleLOM adapter that integrates seamlessly with the server’s architecture and form factor.

FlexibleLOM technology uses a custom implementation of the PCIe 3.0 x 8 interfaces, but does not require any additional CPU resources over standard LOM architecture and does not occupy a regular PCIe slot. FlexibleLOM technology maintains the full functionality of LOM architecture including Wake-on-LAN and thermal management. You can also view the status of a FlexibleLOM on the System Insight Display.

FlexibleLOM technology offers customers of HP ProLiant SL6500 Gen8 scalable servers the flexibility to choose bandwidth speed and fabric type. HP ProLiant SL6500 Gen8 scalable servers can be preconfigured with the FlexibleLOM card that best meets your network requirements or, if you prefer, no card at all. The HP ProLiant SL Gen8 server’s FlexibleLOM configuration is field-upgradeable should connectivity technologies or your network requirements change in the future, thus protecting your investment.

Available FlexibleLOMs include ConnectX3 InfiniBand for ultra-high bandwidth, low latency, and excellent scalability. HP supports the latest 56 Gb/s FDR InfiniBand and 40 Gb/s QDR interconnect along with 40GbE and multiple (up to four) 10GbE connections.

HP and NVIDIA

Since the 1990s, HP and NVIDIA have shared a partnership, developing innovative and powerful workstations for graphic applications. As part of HP's Unified Cluster Portfolio, HP applied cluster technology to visualization and leveraged NVIDIA GPUs, enabling large-scale displays with extremely high resolution. Our shared expertise in HPC and GPUs fostered early collaboration using GPUs for computation, which became a focus area for HP's HPC innovation. HP introduced NVIDIA Tesla GP-enabled servers in 2008, and today ships 14 different GPU-enabled servers. Our customers deploy these servers in a variety of industries, with proven success in oil and gas, financial modeling, life sciences and scientific research. The Accelerator Team within HP's HPC organization has been working with NVIDIA for more than seven years, and HP Labs is one of the newly named Compute Unified Device Architecture (CUDA) research centers.

NVIDIA Tesla GPUs with Kepler architecture

The launch of NVIDIA Fermi GPU in 2009 ushered in a new era in the HPC industry based on a heterogeneous computing model where CPUs and GPUs work together to solve computationally-intensive workloads. With the new NVIDIA Kepler™ architecture, the bar in the HPC industry has been raised again.

Comprised of 7.1 billion transistors, Kepler is created to address the most daunting challenges in HPC. Kepler is designed from the ground up to maximize computational performance with superior power efficiency. The architecture has innovations that make heterogeneous computing significantly easier, applicable to a broader set of applications, and more accessible.

Kepler's computational capabilities include teraflops of integer, single and double precision performance, and high bandwidth memory. The series of Kepler GPUs include Tesla K10, Tesla K20, Tesla K20X, and GRID K2. Kepler's performance features are made possible by SMX, Dynamic Parallelism, and Hyper-Q. To read more about these technologies, go to nvidia.com/object/nvidia-kepler.html.

Tesla K10 GPU Accelerator

Optimized for single precision applications, the Tesla K10 includes two ultra-efficient GK104 Kepler GPUs to deliver high throughput. It delivers up to 2x the performance for single precision applications compared to the previous generation Tesla M2090 GPU in the same power envelope. With an aggregate performance of 4.58 teraflop peak single precision and 320 gigabytes per second memory bandwidth for both GPUs put together, the Tesla K10 is optimized for computations in seismic, signal image processing, and video analytics.

Tesla K20 and K20X GPU Accelerators

Designed to be the performance leader in double precision applications and the broader supercomputing market, the Tesla K20 and K20X GPU Accelerators deliver 10x the performance of a single CPU. Tesla K20 and K20X both feature a single GK110 Kepler GPU that includes the Dynamic Parallelism and Hyper-Q features. With more than one teraflop peak double precision performance, these GPU accelerators are ideal for the most aggressive high-performance computing workloads including climate and weather modeling, CFD, CAE, computational physics, biochemistry simulations, and computational finance.

GRID K2 Boards

NVIDIA's Kepler™-based GRID K2 boards are designed to enable rich remote visualization in HPC environments. NVIDIA's patented low-latency remote display technology greatly improves the user experience by reducing the lag that users feel when interacting with remote visualization resources. With this technology, the remote workstation screen is pushed directly to the remoting protocol. NVIDIA GRID boards have an optimized multi-GPU design that helps to maximize user density. GRID K2 boards (Table 2), which include two higher end Kepler GPUs and 8GB of memory, deliver maximum density for users of graphics-intensive applications. GRID GPUs are designed to provide data center-class power efficiency, including the revolutionary new streaming multiprocessor, called "SMX". The result is an innovative, proven solution that delivers revolutionary performance per-watt for the enterprise data center.

Table 2.

	Tesla K10 ¹ /GRID K2	Tesla K20	Tesla K20X
Peak double precision floating point performance	0.19 Tflops	1.17 Tflops	1.31 Tflops
Peak single precision floating point performance	4.58 Tflops	3.52 Tflops	3.95 Tflops
Number of GPUs	2 x GK 104s	1 x GK 110	1 x GK 110
CUDA cores	2 x 1536	2496	2688
Memory size per board (GDDR5)	8 GB	5 GB	6 GB
Memory bandwidth for board (ECC off) ²	320 GB/sec	208 GB/sec	250 GB/sec
GPU Computing Applications	Seismic processing, signal and image processing, video analytics	Seismic processing, computational fluid dynamics, computer-aided engineering, financial computing, computational chemistry, Physics, data analytics, satellite imaging, weather modeling	
Architecture Features	next-generation Streaming Multiprocessor (SMX)	SMX, dynamic parallelism, Hyper-Q	
System	Servers only	Servers and workstations	Servers only

¹Tesla K10 specifications are shown as the aggregate of two GPUs

²With ECC on, 6.25% of the GPU memory is used for ECC. For example, 6 GB total memory yields 5.625 GB of user available memory with ECC.

- Performance
 - 13x double-precision of CPUs
 - IEEE 754-2008 SP and DP floating point
- Flexibility
 - Increased shared memory from 16 KB to 64 KB
 - L1 and L2 caches
 - ECC on all internal and external memories
 - Enable up to 1 Terabyte of GPU address space
 - High-speed GDDR5 memory interface
- Usability
 - Multiple simultaneous tasks on GPU
 - 10x faster atomic operations
 - C++ support
 - System calls; print support

For complete specifications on the SL6500-compatible NVIDIA GPUs, see “NVIDIA Tesla GPU Modules for HP ProLiant Servers” at: hp.com/products/quickspecs/13743_na/13743_na.pdf

HP HPC management tools

A critical component of any GPU-enabled HPC environment is management. To ease the burden of managing tens of thousands of compute nodes, both CPUs and GPUs, HP offers management tools using the iLO Management Engine as the foundation for reporting system health and status. The HP Insight Cluster Management Utility (Insight CMU) software and the SL-APM are two of the management tools you can use with the ProLiant SL6500 Scalable System. This mature, well-tested GPU-aware cluster management software continues to evolve as each new architecture and operating system is released.

You're providing the highest possible performance in a general purpose HPC data center when you integrate NVIDIA Tesla GPUs into HP ProLiant SL250s Gen8 or SL270 Gen8 servers in an HP Cluster Platform 3000SL and manage those servers with HP management software such as HP Insight CMU and SL-APM. All aspects of this solution are completely supported by HP with a full range of service options.

HP Insight Cluster Management Utility

Supporting all HP Linux-based environments and systems, HP Insight CMU is designed as a light-weight, flexible management system with an intuitive graphical interface that enables you to see your entire cluster. Depending on your needs, you can manage multiple clusters all at once or one at a time.

This "cluster friendly" solution allows you to measure numerous characteristics of the server environment, including memory and rate of I/O reads and writes for each server. It also measures and can set alerts for temperature, fan speeds and hardware health metrics, including GPU metrics. Insight CMU allows you to perform operations on multiple servers, including starting them up and shutting them down. In addition, you can install the OS on one or 1000 servers, all from scratch, in less than an hour. And today, all these valuable cluster management capabilities are available with the new HP ProLiant SL6500 Gen8 servers.

HP SL Advanced Power Manager

The HP SL Advanced Power Manager (SL-APM) is a rack level solution for the SL6500 Gen8 ProLiant Scalable System which enables advanced power management of SL6500s installed in racks using a Command Line Interface (CLI) or scripting over SSH, Telnet or a serial port. Each SL-APM is capable of supporting up to 80 total SL6500 Scalable Systems, and uses HP's SL Rack Dynamic Power Capping capability to monitor and manage power consumption.

SL Rack Dynamic Power Capping gives you the ability to cap the power utilization for SL6500 systems. Using this capability, you can set a power limit, or cap, for each individual SL6500 system that the SL-APM module manages. The SL6500's chassis controller then controls overall power usage for the SL6500 by managing the power utilization of the components in each server node, including that of the fans and the main processors as well as the Xenon Phi coprocessors.

Because it is dynamic power capping, the SL-APM dynamically reallocates the power caps for individual server nodes within the SL6500 based on the nodes current compute workloads. This helps ensure that the enclosure power cap is maintained with as little impact – if any – on overall cluster performance.

SL Rack Dynamic Power Capping with the SL-APM is a powerful tool that allows you to safely provision the appropriate amount of power to Hyperscale rack environments while avoiding possible power outages caused by tripping circuit breakers at peak power. Without dynamic power capping, you would typically be forced to overprovision power by up to 100% of normal consumption to avoid possible outages.

The iLO Management Engine

The iLO Management Engine 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. As a result, many HPC sites minimize monitoring in order to reduce performance impact, which in turn results in a degraded ability to predict server problems before they happen. With Active Health and Agentless Management, all the monitoring is performed on the iLO Management Engine, allowing extensive monitoring without impacting performance. This gives you significantly more information on the state of the system, information letting you predict problems and analyze problems that might be occurring across the cluster.

iLO Agentless Management doesn't use the system CPU, running on the iLO processor instead, and isn't dependent on the operating system. iLO supports a complete separation of system management and data processing, not just on the LAN connections, but also within the system itself. With the iLO Management Engine in ProLiant Gen8 servers, the base hardware monitoring and alerting capability is built into the system and starts working the moment that a power cord is connected to the server.

The Active Health System is an always on acting as a change tracking and run-time diagnostic telemetry system. A rich set of information is recorded from server initialization through its eventual retirement. This capability is built into iLO and requires no setup, no configuration and no SW to install. The SL6500 Gen8 ProLiant Scalable System servers have 4 gigabytes of FLASH storage. We can save the hardware history of the server and up to 2 years of daily runtime telemetry.

Agentless Management, Active Health System, and Embedded Remote Support are game-changing features for iLO 4. All are iLO Standard features included with ProLiant Gen 8 servers. These features do not require additional cost or licensing.

Part of the HP Unified Cluster Portfolio

The HP SL6500 Gen8 Scalable System is available as part of the HP Unified Cluster Portfolio (UCP), a comprehensive suite of solutions designed to help you drive a competitive edge while balancing the pressure on your budget. This innovative, modular package of hardware, software and services is the perfect match for all your scalable computation and data management needs. Serving as the foundation of the HP UCP, HP Cluster Platforms provide a choice of processors, operating systems and interconnects, and includes support for the new SL6500 Gen8 system servers with NVIDIA Tesla and GRID GPUs.

The HP Cluster Platforms 3000 and 4000 are designed with 1U and 2U rack-mount servers as the basic compute building blocks. Easy to configure and ready to run, these platforms provide excellent return on your investment. Able to handle massive scale-out requirements, the newest offerings are the HP Cluster Platform 3000SL and 4000SL, built on an efficient multi-mode chassis that lowers system costs and reduces power consumption. All servers contain multi-core processors to deliver the highest performance possible for your HPC applications. Interconnect choices include Gigabit Ethernet or InfiniBand (QDR or DDR).

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

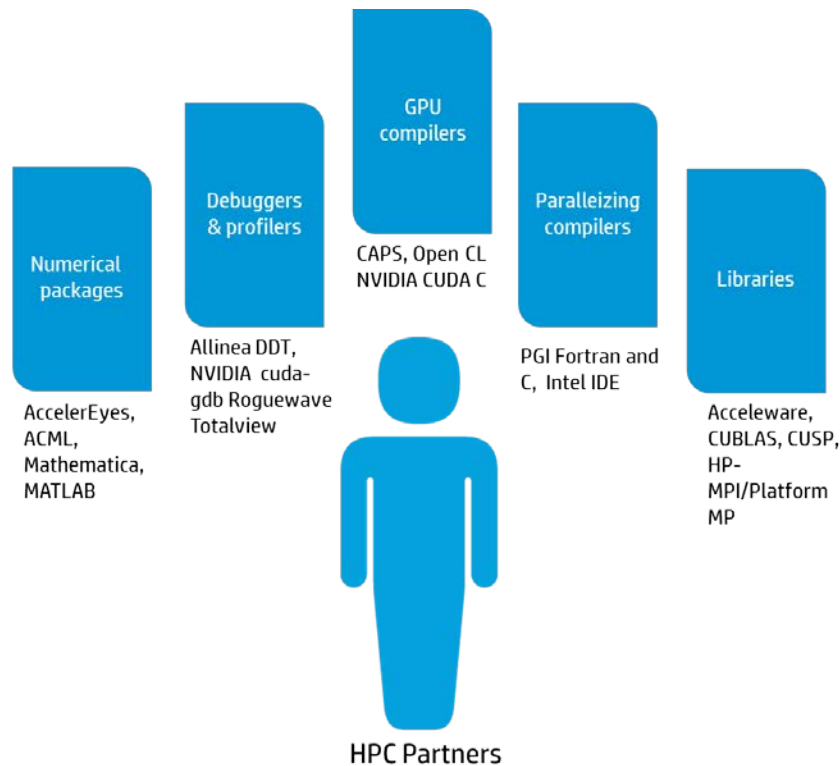
Common deployment worldwide means common serviceability, making HP clusters easy to support and repair. You have a choice of operating systems (Linux or Microsoft® Windows®), cluster management and job schedulers. You also have access to Adaptive Moab HPC Suite, or the Adaptive Workload Optimization Pack (Intel-MPI, Adaptive Moab) as well as Altair PBS® Professional and other third party software.

Your choices for the nodes in your cluster include HP DL and BL server blades, and now next generation SL systems, including the HP ProLiant SL6500 Gen8 system servers with the option to include 1 to 8 Tesla GPUs or up to four GRID K2 boards. You can read more about the HP Unified Cluster Portfolio at: hp.com/HPC/cache/275420-0-0-121.html

HPC partners

The HP HPC partner ecosystem continues to expand offering a growing list of products from industry-leading independent software vendors (Figure 3), making it easier for you to deploy and use GPU-enabled systems:

Figure 3.



A portfolio of partners providing an ecosystem for solutions is critical to the success of any product. To drive the success of our GPU-enabled systems, HP works with a series of partners in three primary areas:

- GPU-enabled applications – industry-specific solutions from numerous partners (described below)
- GPU development environments and tools – to simplify reliable code development
 - PGI compilers and tools – including OpenMP, OpenACC directives, PGPROF, GPU cluster profiler, and PGDBG MPI/OpenMP cluster debugger.
 - Allinea tools - Allinea MAP shows you where the computational bottlenecks in your code are and which would benefit from offloading to CUDA-capable accelerators; Allinea DDT parallel debugger: shows you exactly how your code behaves on the CUDA hardware and detects and helps fix bugs
 - RogueWave – TotalView - Comprehensive tool for verifying, debugging and optimizing complex applications
 - CAPS - GPU development environment and training
 - NVIDIA - CUDA GPU development language
 - Altair PBS Professional - Portable Batch System (PBS) technology helps customers optimize the utilization of enterprise computing environments by intelligently aggregating and scheduling computational resources (workload scheduling).
 - Adaptive Computing - The Adaptive Workload Optimization Pack provides everything to quickly and easily complete an HPC cluster in one integrated package. It provides all the capabilities to optimize workload performance including comprehensive workload scheduling and management, industry-leading messaging passing interface (MPI), and seamless integration with HP Insight Cluster Management Utility (CMU) for reduced administration and improved control.
 - Platform LSF - Fast, scalable service-oriented architecture grid computing middleware; delivers speed-to-value through unparalleled application performance built on top of a utility-based infrastructure-sharing platform, a perfect complement to HP Converged Infrastructure.

- Systems and libraries – the operating environment and common routines upon which applications are built for various industries
 - NVIDIA libraries - Included in the CUDA Toolkit 3.2; also includes supporting drivers, Software Developer’s kit code samples and documentation; available to all GPU computing registered developers
 - Numerical Algorithms Group (NAG) libraries - Offers 1600+ tried and tested routines that are both flexible and portable; remains at the core of thousands of programs and applications spanning the globe
 - MathWorks - MATLAB is the language of technical computing used across several industries with native support for GPUs on your desktop or on clusters; MATLAB users can take advantage of CUDA-enabled devices through built-in GPU accelerated functions or by connecting their CUDA kernels to MATLAB.
 - Wolfram Research - Mathematica routines and libraries for GPUs for various industries
 - RogueWare - IMSL numerical libraries
 - Acceleware - ArrayFire CUDA GPU libraries: library of CUDA and OpenCL functions
 - Acceleware - Provides Parallel Computing software solutions, consulting and training for GPUs to the Oil & Gas and Computer-Aided Engineering markets.

The HP portfolio of partners is constantly evolving. The initial set of partners adapting their products to support GPUs, their focus area and their products are listed below.

Summary

Even if GPUs are beyond the scope of your current computing environment, the HP SL6500 Gen8 Scalable System is still the right choice for HPC, with HP Gen8 systems; HP delivers high performance innovation at any scale

- Scalable performance: New levels of performance that let you speed advancements, with a converged infrastructure that is purpose-built for scale.
- Maximum efficiency: Optimize the performance footprint for every dollar, watt, and square foot with the world’s most efficient systems.
- Instant-On Agility: Deploy easily, adapt quickly to change, and improve quality of service

Contact your HP representative or partner today to learn more about the current generation of HP GPU-enabled servers. You can also visit hp.com/go/hpc to find out more about HP high-performance servers, and how clusters of HP ProLiant SL6500 Gen8 system servers can super-scale your computing environment.

Resources

HP ProLiant SL6500 Scalable System product page
hp.com/wwwpc/us/en/sm/WF04a/15351-15351-3896136-3896139-4236125.html

HP ProLiant s6500 Chassis QuickSpecs
hp.com/products/quickspecs/14188_na/14188_na.pdf

NVIDIA Tesla GPU Modules for HP ProLiant Servers QuickSpecs
hp.com/products/quickspecs/13743_div/13743_div.pdf

NVIDIA Tesla GPU Modules for HP ProLiant Servers
hp.com/products/quickspecs/13743_na/13743_na.pdf

HP InfiniBand Options for HP ProLiant and Integrity Servers
hp.com/products/quickspecs/13078_na/13078_na.pdf

HP Insight Cluster Management Utility
hp.com/go/cmu

iLO product information and user guide
hp.com/go/ilo

ProLiant server information
hp.com/go/ProLiant

Industry Standard Server Technology Papers
hp.com/servers/technology

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 Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are U.S. registered trademarks of the Microsoft group of companies.

