

ACCELERATE

high-performance computing.

*Clustered HP ProLiant SL390s
servers with NVIDIA GPUs*

Solution brief



Reach new levels of innovation and competitiveness with our second generation of GPU-enabled servers.

The trend is clear—high-performance computing (HPC) is now mainstream and is pushing the boundaries of innovation and business performance like never before. Businesses are using HPC technology for faster innovation and competitive advantage. HP brings HPC within the reach of more organizations by offering complete supercomputer clusters of purpose-built ProLiant servers with x86 processors. Our approach differs from competitors that stipulate long-term commitments to costly, proprietary supercomputers. Our industry-standard ProLiant servers accommodate high-end graphic cards and accelerators and help raise performance and drive down costs in HPC visualization and computation environments. We jointly engineer and test our systems with products from industry-leading accelerator vendors and independent software vendors (ISVs) for a diverse range of HPC environments, including oil and gas, biosciences, financial services, government, automotive, and aerospace.

Accelerate business outcomes with GPU-enabled computing

What HPC GPU-enabled computing offers your business is lightning fast time to discovery, deeper insights, and ability for smarter decision-making resulting in faster time to market at a fraction of cost as that of traditional systems.¹ One of the key emerging technologies in HPC are Accelerators—coprocessors designed for heterogeneous processing that provide dramatic computation speeds for compute intensive applications. Leading the race in accelerator technology are graphics processing units (GPUs). GPU-enabled systems have dramatically skyrocketed performance—giving greater performance/watt, better price/performance, and providing significantly more compute cycles. Based on the parallel processing paradigm, GPUs offload computing from CPUs, accelerating application performance by leaps and bounds.

Prime GPU performance with the SL390s

HP is at the forefront of HPC innovations, building GPU-enabled platforms from the ground up on our industry-leading ProLiant range. These server systems are qualified with the requisite software and demonstrate benchmarked performance, when measured end to end. The second generation of GPU-enabled hybrid servers, HP ProLiant SL390s G7 servers are purpose built to be the premier platform for GPU-computing. They are green and they are low on energy, footprint, and operational costs. The June 2011 Green500 list rates the TSUBAME 2.0 built with SL390s at Tokyo Institute of Technology as the “world’s greenest production supercomputer.”²

Meet the HP ProLiant SL390s servers

The highly scalable ProLiant SL390s servers take GPU-enhanced cluster computing to the next level. These servers are designed for the extreme-scale dynamic workloads HPC computing demands, and can achieve high server and rack density at outstanding power efficiency. Depending upon the application, a SL390s GPU-enabled system offers 5-10x,¹ and potentially more, the performance of traditional servers—meaning faster and affordable resolution to your business challenges. Making management of tens of thousands of compute nodes simple is the powerful HP Cluster Management Utility (CMU) that provides a “single pane of glass” approach to manage clusters and GPUs within the cluster. The HP Unified Cluster Portfolio (UCP) provides a comprehensive Converged Infrastructure approach to HPC solutions and is designed to accelerate innovation.

The ProLiant SL390s combines key learning and best practices from our previous generations of HP ProLiant line of servers, including density, efficiency, and shared fans and power. The result is a series of massively scalable systems that are designed for the hyper-scale computing demands of our customers—purpose built for scale and performance with integrated system and GPU management, easy accessibility, and efficient power and cooling.

Purpose built for scale

The new ProLiant SL390s G7 servers offer all the richness of the HP ProLiant heritage in a massively scale-out, second-generation, GPU-enabled server.

Features of the HP ProLiant SL390s models:

- Built on the HP ProLiant s6500 chassis
- HP ProLiant SL390s G7 Server trays
 - 1U half-width, built for CPU density
 - 2U half-width, supports up to three NVIDIA Tesla GPUs
 - 4U half-width, supports up to eight NVIDIA Tesla GPUs
- Half width for maximum density—eight servers in 4U with the 1U server, four servers in 4U with the 2U server, and two servers in 4U with the 4U server
- Integrated 10GbE/InfiniBand (IB) and optional dual IB rail
- HP Integrated Lights Out-3 capabilities
- Front serviceability
- Dedicated or shared (server dependant) PCIe Gen2 x16 lanes to the GPUs—to exploit GPU capability fully

Options for the HP ProLiant SL390s GPU-enabled hybrid servers include:

- Increased density of GPU-to-CPU cores
- Hot-swap drives
- Robust management: advanced power management and GPU management and monitoring

¹ HP Internal: According to the June 2011 Green500 list, Jaguar is #3 with an R-peak of 2.3 Petaflops and the GPU-enabled TSUBAME 2.0 is #5 with an R-peak of 2.28 Petaflops. However, while Jaguar uses 18,688 servers. The TSUBAME 2.0 offers similar performance with only 1,408 SL390s servers—almost 17,000 fewer servers. www.green500.org

² Source: [Creating an environment-friendly supercomputer with the HP ProLiant SL Scalable System](#), Tokyo Institute of Technology, author: HP, year 2011

Why HP and NVIDIA?

Since the 1990s, HP and NVIDIA have shared a deep partnership, developing innovative and powerful workstations for graphic applications. As part of its UCP, 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 by using GPUs for computation, which became a focus area for innovations in HPC. Our customers deploy these servers in applications for oil & gas, financial modeling, life sciences, and scientific research. The Accelerator team within our HPC organization at HP has been working with NVIDIA for more than four years. Also, HP Labs is one of the newly named CUDA Research Centers.

Why use GPUs?

Once considered difficult to program and use, GPUs are moving to the mainstream. Today, an emerging software development environment—complete with platforms, compilers, debuggers, and libraries—is available from NVIDIA and partners. Leading software developers are building applications using the emerging GPU development environment and the NVIDIA programming environment—called CUDA. Solution providers are finding it increasingly easy to build applications using these tools. GPUs are here to stay. Applications are complete or in process for all the major HPC markets—from oil & gas to life sciences.

GPU-enabled systems have the potential to offer:

- Significant compute, power utilization, and performance benefits
- Hundreds of functional units executing in parallel
- Ability to speed up applications by 2x, 10x, 30x, or up to 100x in roughly the same server space
- Excellent support for the highly parallel, compute-intensive applications typically used in HPC environments

Unified management, built in

A critical component of any GPU-enabled HPC environment is management. To ease the task of managing tens of thousands of compute nodes—both CPUs and GPUs—HP offers the Cluster Management Utility (CMU). Entering its tenth year of development and used in clustered environments, this GPU-aware software continues to evolve with each new release.

Supporting all HP Linux-based environments and systems, CMU is designed as a light-weight, flexible management system with an intuitive graphical interface that enables you to see your entire cluster. CMU simplifies management, monitoring, and provisioning across a cluster or a simple group of nodes.

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—and now GPU metrics such as GPU utilization and temperature. In addition, you can install the OS on one or 1000 servers, all from scratch, in less than two hours.³

The HP difference


- **HP ProLiant SL390s servers:** With both dedicated x16 lanes for each GPU in the 2U and shared x16 lanes in the 4U variant, the SL390 is designed to support GPUs from the ground up with our manufacturing, worldwide support, and integrated management. Compare this with other vendors who treat GPUs as mere server add-ons.

- **HP Cluster Platform 3000SL/4000SL:** The market leader in HPC clusters, we offer supported and integrated HP Cluster Platforms, available worldwide for rapid and confident deployment. ProLiant SL390s servers with GPU options are standard for the HP Cluster Platform 3000SL and 4000SL, and available with your choice of OS and cluster tools.
- **Native connectivity:** The HP ProLiant SL390s motherboard includes native IB and 10GbE capabilities. You can immediately connect the HP ProLiant SL390s Server to any switch. The integrated technology approach lowers price, offers greater reliability, and predictability. Also, the SL390 provides an additional PCIe slot for custom I/O requirements.
- **Greater choice:** Choice of OS (Linux or Microsoft® Windows®), cluster management, and job schedulers, as well as HP-MPI, Linux Value Pack (HP-MPI, LSF, and UPC with SHMEM)
- **Rich ecosystem:** With its development team and partners, HP fosters GPU-enabled development environments and tools, cluster management, and systems and libraries, creating the operating environment, common routines, and solution components for industry-specific solutions.

³ HP Internal: In a customer site implemented with HP CMU, the production system image was cloned to 800 HP BladeSystem servers in 24 minutes.

Accelerate benefits to business with GPU-enabled HPC

Whether your organization focuses on gene sequencing, protein docking, risk analysis, or interactive video, HP ProLiant SL390s servers with NVIDIA Tesla GPUs can accelerate business outcomes that matter. To learn more, visit www.hp.com/go/hpc and www.hp.com/go/accelerators today.



Get connected

www.hp.com/go/getconnected

Get the insider view on tech trends, alerts, and
HP solutions for better business outcomes

© Copyright 2010–2011 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.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA2-8923ENW, Created November 2010; Updated August 2011, Rev. 1

