

# ACCELERATE

high-performance computing.

## Clustered HP ProLiant SL390s servers with NVIDIA GPUs

Reach new levels of innovation and competitiveness with HP's second generation of GPU-enabled servers

Solution brief

### The next technology wave

As a company that relies on high performance computing (HPC) to drive your organization, you've witnessed HPC evolve from very costly, proprietary, UNIX®-based systems to affordable, x86, Linux, cluster-based systems.

Today, the next technology wave has arrived—catapulting HPC sites to super-scale performance while delivering outstanding efficiency and affordability. That wave is GPU (graphical processing unit) accelerated computing. And the new HP ProLiant SL390s G7 Server—HP's second generation of GPU-enabled systems—has been specially designed to be the industry's best platform for GPU computing.

The perfect solution for clusters with GPU workloads, highly scalable HP ProLiant SL390s servers enable you to achieve density at server and rack level while also providing outstanding power efficiency. With up to 150x the performance power of traditional servers, the HP ProLiant SL390s servers will help you resolve bigger problems in less time more affordably.

### Meet the HP ProLiant SL390s servers

In designing the HP ProLiant SL390s servers, we merged key learnings from our BL line—including density, efficiency, and paired fans and power—with best practices from our first-generation HP ProLiant SL

servers. The result is a series of massively scalable, “skinless” systems (slide-out trays housed in a chassis) that offer integrated application and power management, easy accessibility, efficient power and cooling, and more.

While HP was actively designing the second generation of the HP ProLiant SL line, NVIDIA ([www.nvidia.com](http://www.nvidia.com)) was hard at work advancing its Tesla GPU accelerators, now in their third generation. HP recognized the considerable potential of this innovative technology, 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, which included incorporating capabilities into the new HP ProLiant SL390s servers to fully exploit the power of NVIDIA Tesla GPU accelerators. As the first vendor to take this leap in innovation, HP elevated NVIDIA GPU accelerators from a specialized function to a mainstreamed solution.

Offering I/O expandability and networking, HP ProLiant SL390s servers with up to three NVIDIA Tesla GPUs can compact 96 GPUs into a single, industry-standard rack, enabling this “skinless” solution to deliver excellent compute density, power efficiency and affordability.



## 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 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. Our customers deploy these servers in applications for Oil & Gas, Financial Modeling, Life Sciences and Scientific Research. The Accelerator Team within HP's HPC organization has been working with NVIDIA for more than four years, and 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

## Purpose-built for scale

New HP ProLiant SL390s G7 systems offer all the richness of the HP ProLiant heritage in a massively scale-out, second-generation, GPU-enabled server. The 1U server allows more CPU density, and the 2U servers support up to three integrated NVIDIA Tesla GPUs. Both HP ProLiant SL390s models offer:

- Integrated 10GbE/IB and optional dual IB rail
- HP Integrated Lights Out-3 capabilities
- Half-width for maximum density – 8 servers in 4U with the 1U server, 4 servers in 4U with the 2U server
- Front serviceability

Options for the HP ProLiant SL390s 2U servers include integrated NVIDIA GPUs and hot swap drives. Each GPU has a dedicated x16 lane to ensure you can fully exploit the GPU capability.

## Unified management, built in

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 the Cluster Management Utility (CMU). Entering its tenth year of development and use in clustered environments, this GPU-aware software continues to evolve as each new architecture is released.

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 operating system on one or 1000 servers, all from scratch, in less than two hours.

## The HP difference

- **HP ProLiant SL390s servers** – offering higher density of GPUs to CPU cores that means more computing power; up to three times the GPU density of other GPU-enabled systems; dedicated x16 lanes for each GPU; The 2U version supports GPUs from the ground up with HP manufacturing, worldwide support and integrated management, compared to other vendors who treat GPUs as a server add-on.
- **HP Cluster Platform 3000SL** – The market leader in HPC clusters, HP offers supported and integrated HP Cluster Platforms, available worldwide for rapid and confident deployment. HP ProLiant SL390s servers with GPU options are standard for the HP Cluster Platform 3000SL, and available with your choice of OS and cluster tools.
- **Native connectivity** – The HP ProLiant SL390s motherboard includes native InfiniBand and 10GbE capabilities. You can immediately connect the HP ProLiant SL390s Server to any switch; this technology lowers price, offers greater predictability and provides an additional slot.
- **Greater choice** – choice of operating system (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.

## Ready to accelerate?

Whether your organization focuses on gene sequencing and protein docking, risk analysis or interactive video, HP ProLiant SL390s servers with NVIDIA Tesla GPUs can accelerate business outcomes that matter.

Contact your HP representative today. To learn more, please visit [www.hp.com/go/hpc](http://www.hp.com/go/hpc) and [www.hp.com/go/accelerators](http://www.hp.com/go/accelerators).

Share with colleagues    



## Get connected

[www.hp.com/go/getconnected](http://www.hp.com/go/getconnected)

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

© Copyright 2010 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. UNIX is a registered trademark of The Open Group.

4AA2-8923ENW, Created November 2010

This is an HP Indigo digital print.

