High-Performance Computing in Action

In This Paper

- Today's supercomputers adhere to industry standards, with the majority of combinations being x86-based servers running Linux
- Successful HPC solutions are more about problem solving than processing power, and as such they are closely aligned to operations
- HPC fuels innovation in many industries and plays an increasingly key role in ongoing business concerns

ServerWatch™
Executive Brief

High performance computing, sometimes referred to as supercomputing, has been finding its way into the mainstream with newfound popularity in recent years. Once the domain of high-end research facilities and large government agencies, the rise in Big Data and specialized applications that require immense amounts of mathematical calculations are now fueling its growth.

High performance computing (HPC) is made possible by supercomputers and the software that runs on those computers. A key tenet of HPC is the parallel processing algorithms and software containing them. Such programs can be divided into little pieces so that each piece is executed simultaneously by separate processors.

Supercomputing has historically been the apex for cutting-edge computing, and it is arguably as old as computing itself. The earliest computers were, by their very nature, supercomputers. There is, however, no hard and fast rule for what is considered HPC, as the power for what constitutes an HPC system is relative to other compute devices in play at any given time. In other words, HPC is a moving target.

Not surprisingly, HPC today is vastly different from HPC of five decades ago for reasons far beyond size and speed. The earliest supercomputers were proprietary closed systems, tied closely to the vendors that built them and developed the software that ran on them. These earliest supercomputers were built for specific applications and purposes. They were expensive and unscalable, typically requiring significant capital investment. They also occupied an extremely large footprint, which meant significant space to house and maintain them was required. They also often required an investment in human capital in the form of specialization and training.

Ironically, these behemoth installs had less compute power than the smartphone you can easily slip into your pocket today. A typical desktop or laptop today could almost be considered HPC by older definitions.

Today's supercomputers, in contrast, are frequently clusters of commodity hardware with high speeds that are measured in Teraflops, with Exaflop performance expected to be achieved by 2018. (A FLOP is the number of FLoating Point Operations performed Per Second; an Exaflop is 10^{18} or 1,000,000,000,000,000,000 FLOPs.)

As recently as one decade ago, HPC was a clearly defined niche. Today, HPC is pervasive and in many cases requires little specialization beyond industry acumen to accompany standard IT knowledge.

This is because, the rare exception aside, today's supercomputers adhere to industry standards, for both hardware and software. The most common combination is x86-based servers running Linux. Developments, such as multicore processors from vendors like Intel®, and graphical processing unit (GPU)-enabled processors that help with data visualization along with enhanced memory, have removed the low-end performance stigma previously associated with x86 architectures. This has resulted in some pretty formidable machines, especially when clustered together. This sea change is even more evident when reviewing the twice yearly Top500 Supercomputing

list, which each time contains a greater number of x86-based systems running standard operating systems.[1]

Standardization on the x86 architecture meant enterprises were no longer confined to proprietary operating systems. Both Windows and Linux benefitted from this, but Linux has proven to be the bigger winner here. Its scalability and customizability makes it ideal for HPC installs, and Linux file systems facilitate the desired throughput.

Whether clustered or scaled out, HPC solutions typically involve more than one server, which means a variety of networking gear is required. Adhering to industry standards around hardware offers the option to use not only less -expensive processors, but also other industry-standard and, in some cases commodity, networking gear and interconnects.

Hence, no longer is HPC an expensive undertaking. Today, an HPC system is as likely to be a blade server as a mainframe. This alone takes it out of the province of governmental agencies and the scientific university communities and into the enterprise world. Cloud computing takes it a level further, making HPC an option for smaller organizations that cannot afford to run HPC on site.

A growing number of cloud providers are offering HPC services in the cloud. For organizations with an intermittent need for HPC power, for which infrastructure constraints and capital budgets made HPC a nonstarter, the cloud presents numerous possibilities.

In a cloud computing model, resources are available as needed and costs are related to usage. IT Business Edge blogger Arthur Cole cites Amazon's EC2 cloud, which can spin up to 50,000 cores backed by a highly parallel networking architecture built around a 10 GbE fabric. The company says it is adequate for 99 percent of today's HPC jobs. Typical deployments consist of a pair of eight-core Intel® Xeon E5-2670 processors, 60.5 GB of memory and 3.3 TB of storage, although the company can deliver paired quad-core Intel Nehalem instances as well.[2]

"HPC is constantly evolving due to consistent improvements in performance of processors, systems, interconnects, and, most recently, analytics software."

HPC is constantly evolving due to consistent improvements in performance of processors, systems, interconnects, and, most recently, analytics software. Ultimately, however, the driving force is enterprise need.

HPC Today

How is HPC being defined these days, and how do these machines differentiate themselves from standard servers? What sets, for example, the HP ProLiant Scalable System, or SL, line apart from HP's, in some cases equally powerful, other PoLiant servers?

HPC is all about high throughput and parallel operations. Not every big task requires HPC, but tasks that can be done quicker when performed in parallel, benefit from it. This was initially found to be particularly useful for scientific data modeling.

Today, HPC is a conduit for the analytics that turn ever growing mountains of data into information. Nearly every industry reliant on Big Data could stand to benefit from HPC. Big Data analysis has been a common theme in HPC for a long time. Examples include analyzing weather forecast predictions for climate analysis or modeling car crash simulations and comparing them to crashing real cars.

HPC solutions, therefore, are about more than processing power. They are about problem solving. Ideally, an HPC solution will be closely

aligned to operations such that it reduces IT management time and costs, along with decreasing power and cooling expenses (typically the biggest data center cost). Add in general affordability, and you have a trifecta that would be difficult for any enterprise to turn down.

HP's HPC solutions achieve this trifecta. From the data center to desktop, the company has a long history in the HPC market. Its wide array of HPC solutions is a key factor in its status as a leader in this space. HP has consistently strong showing on the twice yearly seminal Top500 supercomputing list and is a dominant player as measured by the number of units shipped. HP also dominates in terms of HPC solutions chosen in the oil and gas, aerospace and automotive industries.

This is due, no doubt, to HP's spectrum of solutions that meet these needs. This is evident in the HP ProLiant Scalable Systems line, in particular the SL200 Gen8 series, an advanced Intel®-based server family that consists of three 2-socket server trays, each offering unique form factors and benefits: the SL230s, 1U half-width server tray, is designed for compute density and featuring an optional integrated PCIe IO Accelerator; the SL250s, a 2U half-width trays, is designed for GPU density, with extra space for support for up to three GPU slots or additional storage; and the SL270s, a 4U half-width tray, is also designed for GPU density, with support for up to eight GPU slots.

As its name implies, the HP ProLiant Scalable Systems line was designed to provide the scale service providers seek for HPC and Big Data. The SL series integrates computing, storage, management software and networking, as well as a power and cooling infrastructure, intended for an HPC environment. And it does so in a way that offers rapid deployment, greater agility and lower operational costs.

Each of the three SL servers offers a different value proposition. The SL230s and SL250s feature high levels of configurability for high performance and scale-out deployments and deliver a lower total cost of ownership, while the SL270s reduces solution costs, maximizes power efficiency and offers easier serviceability.

With HPC comes amped-up storage needs, particularly around Big Data. To address them, HP offers the SL4540 Gen8, a server built specifically to solve the traditional structured database problems that crop up when dealing with Big Data. The SL4540 features a highly efficient converged design that delivers the ideal combination of capacity and performance in the least amount of space and at lower cost, with high reliability and manageability.

The HP ProLiant SL4540 Gen8 comes in three flexible compute nodes, all featuring the Intel® Xeon® E5-2400 product family, contained in a 4.3U form factor chassis that allows for shared cooling, power and management, as opposed to siloed architectures with individual chassis.

Enterprises seeking a smaller footprint for their HPC will find their needs met with the HP ProLiant server

blades. Two standouts are the Intel®-processor-powered BL460c Gen 8 and BL660 Gen8. The BL460c is a dual-socket server blade engineered for unprecedented performance, enhanced flexibility and simplified management.

The BL460c packs in more performance with a 33 percent increase in memory DIMM count, Intel® Xeon® E5-2600 processors with added support for 135W processors, faster I/O slots and an enhanced Smart Array Controller that now ships with 512MB Flash Back Write Cache standard.

The HP ProLiant BL660c Gen8 Server series, meanwhile, is ideal for virtualization, database, business processing, and general 4P data-intensive applications where fine-tuning of data center space and price/performance is paramount. Powered by the high-density and cost effective Intel® Xeon® E5-4600 processors, the BL660c Gen8 Server series redefines density-optimized 4-socket blade technology without compromising on performance, scalability, and expandability.

Hardware enhancements include the HP SmartDrives, Smart Socket guide, SmartMemory, FlexibleLOM technology, "3D Sea of Sensors" and the embedded Smart Array P220i RAID controller.

Whatever servers an enterprise chooses, the HP Unified Cluster Portfolio (UCP) makes it easy to tie them together into a cohesive HPC system.

"Today, the window between theory and practicality grows ever narrower as HPC's relevance to ongoing business concerns grows stronger."

Servers from the SL and BL lines can be rolled upward into the HP Unified Cluster Portfolio (UCP). The UCP is designed for enterprises looking for to bring together multiple operating systems, and thus provides a common implementation for operating systems, including Microsoft Windows and Linux, and can be used with servers based on industry-standard processors, such as Intel® Xeon.

The heart of UCP with Intel Xeon is the HP Cluster Platforms 3000, which provides a foundation and bridge between servers. This comprehensive portfolio of fully integrated, standards-based HPC solutions and services can be customized for various industries, such as energy.

HP Cluster Platforms 3000 are now available with the ProLiant Gen8 servers featuring HP ProActive Insight architecture, new Intel Xeon processors including the E5-2600 family, and fast FBR InfiniBand.

Enterprises not seeking this level of customizability can take advantage of HP Cluster Platform Express, which is factory assembled and ships ready for deployment with pre-installed software for cluster management. Businesses have a choice of ProLiant DL and BL server and software options for HPC with up to 32 compute nodes in a single cabinet with rack servers and up to 48 blades with a c-Class Blade chassis.

HPC in the Enterprise World

While it's arguable that many of the exciting and groundbreaking efforts of supercomputing have little impact on most business' day-to-day reality in the immediate term, supercomputing drives innovation and is a vital part of the big picture for numerous industries. In the past, developments in supercomputing drove innovation first in science, technology, engineering and mathematics (STEM) fields, and then progressed to other industry segments, with the high-end impacting the low-end and midrange. Today, the window between theory and practicality grows ever narrower as HPC's relevance to ongoing business concerns grows stronger.

Numerous industries have benefited from HPC. HPC has fueled innovation in industries as diverse as energy, manufacturing, financial services

and entertainment. Such innovation changes the playing field for enterprises within these industries and has enabled better products to be developed or analysis to be made and acted on in less time and at lower cost.

The energy industry was one of the earliest adopters of HPC and perhaps best exemplifies the value that HPC brings to the operations. Modeling is a critical component for the oil and gas industry when it comes to seismic exploration and reservoir simulation. HPC facilitates this modeling in a way that cannot be replicated. Even with modeling, however, acquiring new reserves is becoming increasingly costly, particularly as exploration and production extends into deeper water offshore or hostile climates in remote locations onshore. There is a tremendous amount of data to be gathered and analyzed, posing an enormous task for the engineers and geoscientists who must evaluate where and how to drill. As costs around the actual procurement increase, there is a drive to make the other aspects, including the data, technology, people and processes more cost-efficient.

The HP Unified Cluster Portfolio (UCP), in particular, the HP Cluster Platforms 3000, offers a comprehensive portfolio of fully integrated, standards-based HPC solutions and services developed specifically for the oil and gas industry. The HP ProLiant SL270s Server is also ideal for HPC applications in oil and gas as well as life sciences, manufacturing, research and financial services.

MicroSeismic, Inc. is one company in the energy industry that relies on HPC to remain a going concern. MicroSeismic is an oil field service company that offers real-time monitoring and mapping of hydraulic fracture operations for the development of unconventional oil and gas extraction. It provides "passive" seismic monitoring of oil and gas assets with recordings at or near the surface, listening for microseismic energy generated in the oil-and-gas-bearing rock formation during stimulation or production. It then passes its findings along to clients, who use this data to decide where to drill, and how to complete their wells and develop their fields. In some cases, data must be processed in real time so MicroSeismic can show customers the impact of their fracturing jobs as they happen, enabling them to make adjustments to operations at any given stage.

Faced with a backup due to an IT infrastructure that could not handle the demands of processing data in real time, MicroSeismic turned to HP and HP Elite Partner Technical & Scientific Application, Inc. (TSA). Using HP Converged Infrastructure-based technologies, the company was able to increase server management eight-fold, achieve faster data processing for optimizing well stimulation using HP S6500/SL390 and BladeSystem, reduce power consumption 50 percent with HP S6500/SL390 and BladeSystem, and achieve business continuity with data replication across two sites. As significant as these gains were, the business benefits were even more significant — a 20 percent reduction in TCO in server hardware thanks to virtualization, an increase in

revenue due to faster turnaround and improved productivity, an increased ability to service more customers, and a $100,000 savings in equipment attributable to the HP Storage P4000 storage area network.

A MicroSeisimic-like HPC story can be found in nearly any industry. In manufacturing, for example, HPC has had an impact across the board, through the entire process and from the data center all the way to desktop.

HPC plays a key role in product development. The aerospace industry, for example, relies heavily on modeling for product development. It would be cost-prohibitive and virtually impossible to test for every possible circumstance. Companies must therefore rely on simulated models that require enormous compute power. With HPC, engineers simulate conditions inside and outside the aircraft.

Hence, aerodynamics, structure, acoustics and avionics systems are tested to ensure safety, expected functionality and performance, and comfort and as well as adherence to environmental regulations. [3]

HPC has also had a profound impact on the financial services industry. While no actual products are created, financial firms are constantly making decisions based on large volumes of data while creating large volumes of data on their own.

Not only do financial services firms need to transact faster with each passing year to remain viable, but they must also contend with a wider reach. HPC ensures trades are performed with the high-speed, low-latency combination on which firms operating in a global environment now depend. HPC also helps ensure accuracy, reliability and security for these high-speed trades that process data from multiple streams. HPC also comes into play behind the scenes in algorithmically driven analysis. It has the computational chops to sort through data-intensive workloads, generate predictive analysis, and in many cases execute a decision based around its Big Data analysis, often without any human intervention.

HPC is also serious business for the entertainment industry. Digital media relies heavily on HPC, for rendering in gaming and movies. The entertainment industry harnesses the power of HPC to create computer-generated animations to create an animated feature film. No longer must a large pool of animators be on staff to draw every film cell. This process took a long time and was often be extremely costly. Using HPC allows studios to produce more movies in less time at a far lower cost.

HPC is also used to generate or augment special effects and scenery in live action films. This, too, keeps production costs down while enhancing the experience for the viewer.

The Future of HPC

All of these industries and others rely on HPC. Through the years, HPC has become increasingly pervasive, touching many aspects of nearly every industry from the data center to the desktop and finally to the end user. It has changed end user perspectives on what is possible, and moves the bar ever higher on what is expected. As Big Data grows ever bigger, this bar will continue to rise higher and faster.

Enterprises must be prepared to keep up.

HP has been there every step of the way. Through its partnership with the leading hardware and software players in the industry, including Intel®, HP is able to offer customers the latest innovations in processing power and applications from the data center all the way down to the desktop or, in some cases, in the field. ■

[1] http://www.serverwatch.com/trends/article.php/3931791/HP-ProLiant-Scalable-System-Buyers-Guide.htm

[2] http://www.itbusinessedge.com/cm/blogs/cole/pushing-the-cloud-to-hpc-levels/?cs=50512

[3] http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA3-9976ENW.pdf