White Paper

The Business Value of Leading-Edge High Performance Computing: 2019 Update

Steve Conway, Earl Joseph, Robert Sorensen and Alex Norton

March 2019

Sponsored by Hewlett Packard Enterprise and Intel

HYPERION RESEARCH OPINION

High performance computing (HPC) first proved its transformational value serving government and academic researchers in the early 1960s. HPC began penetrating Tier-1 commercial firms in the late 1970s and quickly became a game changer for accelerating innovation, product design and complex business processes. Further expansion into companies of all sizes helped to propel the global HPC market from $2 billion in 1990 to $24 billion in 2017, en route to a Hyperion-forecast $38 billion in 2022.

HPC's value for business, government and academic users will ramp up further as a uniquely powerful resource for Big Data analytics, including machine learning, deep learning and other AI methods. It should be no surprise that businesses and other leading organizations in a wide range of sectors are striving to out-compete their rivals by using HPC to out-compute them—not just by exploiting HPC, but by deploying the highest-capability, highest-capacity HPC systems available in the market place.

HPC problems with the highest value for user organizations, including new Big Data analytics problems, also tend to be the most technically challenging. They typically require leading-edge supercomputers: systems with prodigious memory capacities, processing and communication rates to produce verifiable results in needed timeframes. Leading-edge supercomputers can reduce time-to-value for daunting simulation problems from months to hours, and slash solution times for today's most demanding Big Data problems, including emerging AI, machine and deep learning tasks, from a week or more to as little as 50-150 milliseconds.

Some of today's most difficult problems will migrate over time to less-capable HPC systems or even desktop computers, but there is no shortage of new, even more demanding business and research problems to keep future leading-edge supercomputers busy. Hewlett Packard Enterprise (HPE), the revenue leader in the worldwide HPC market, has long been a powerhouse across all HPC system price bands, from sub-$100,000 Workgroup systems to $500,000-plus Supercomputers. More recently, HPE has begun successfully targeting the global market for leading-edge supercomputers with the company's HPE Apollo and HPE SGI 8600 product lines outfitted with the latest Intel® Xeon® Scalable processors and HPE's system and application acceleration software solutions designed for extreme-scale supercomputers. Hyperion Research believes that HPE is well positioned to become an increasingly important player in this financially rewarding, prestigious portion of the HPC market, for both on-premises deployments and the growing use of HPC in cloud environments.

STRONG PROSPECTS FOR HPC MARKET GROWTH

During the past 15 years, HPC has been one of the fastest-growing IT markets, at times surpassing the annual growth rates of online gaming and flat-panel TVs. In full-year 2017, we reported that worldwide factory revenue for the HPC server market grew to a record $12.3 billion, up from $11.2 billion in 2016 and more than double the 2003 figure of $5.7 billion (see Figure 1). Hyperion Research forecasts that the HPC server market will grow to $19.6 billion in 2022, and that the market for the whole HPC ecosystem (servers, storage, software and support) will exceed $38 billion in that year.

FIGURE 1

Worldwide HPC Server Market Growth ($ Billions)

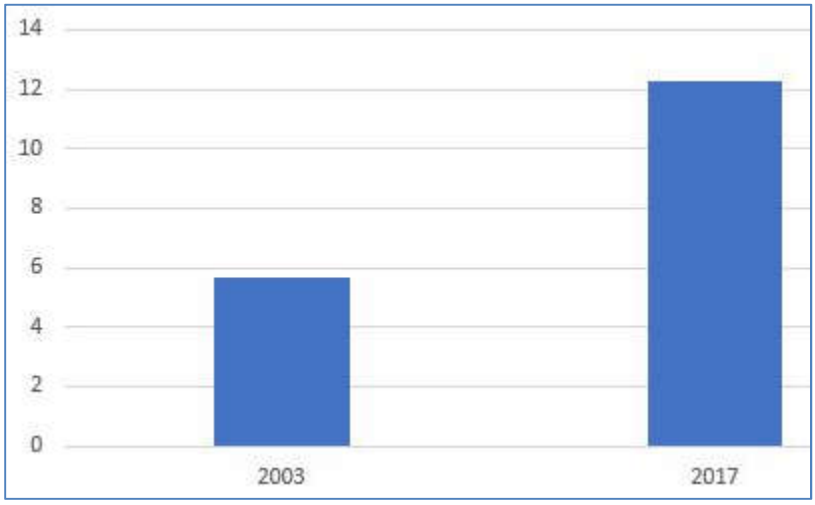

Source: Hyperion Research, 2019

Quantifying the Value of HPC for Businesses

Half of Spending on HPC Systems Is by the Private Sector

Of the $12.3 billion revenue total for HPC servers in 2017, just over half ($6.3 billion, or 51.2%) came from purchases by private sector businesses. This heavy adoption by the private sector runs counter to a popular misconception that HPC remains rooted in its 1960s origins as a specialized market serving only government and academic researchers. In reality, large businesses began adopting HPC in the late 1970s to speed and improve the design of complex products such as aircraft and automobiles. A decade later, following the arrival of affordable entry-level HPC systems, HPC modeling and simulation expanded downstream into smaller companies and other organizations, as well as into workgroups and departments within larger organizations. In 2017, some 88,600 HPC systems were sold worldwide and 81,400 of them (91.9%) were in the sub-$250,000 Workgroup and Departmental price bands.

HPC Provides Large Returns for Businesses

97% of Business Adopters Find HPC Indispensable

In an earlier study we conducted for the Washington, DC-based U.S. Council on Competitiveness, 97% of private sector firms that had adopted HPC said they could no longer compete or survive without it. A 2017 NSF-sponsored Hyperion Research global study for the National Center for Supercomputing Applications confirmed that businesses that exploit HPC benefit from accelerated/improved product design and increased competitiveness.

A $1 HPC Investment on Average Yields $463 in Revenue and $44 in Added Profit

Following a successful 2013 pilot study for the U.S. Department of Energy, for which we developed two macroeconomic models for measuring financial and non-financial ROI associated with government investments in HPC, DOE authorized us to conduct a full-out, multi-year worldwide study to populate the models with data from thousands of real-world, public- and private-sector projects from across the globe. Results from this study-in-progress indicate that, on average, every $1 invested in HPC is associated with $463 in additional revenue and $44 in added net profit or cost savings for businesses that exploit the HPC investments (see Figure 2, and the full data at: www.hpcuserforum.com/ROI). Hyperion Research has completed studies using the same methodology for government agencies in Germany, the Netherlands, Japan, the Republic of Korea and for the National Science Foundation (U.S.). Figure 2

FIGURE 2

ROI Associated with HPC Use, by Market Sector

Q3: Industry	Sum of Total Jobs Added	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $
Academic	14	$9 M	$44 M
Defense		$75 M	$5 M
Financial	602	$834 M	$61 M
Government	42	$1,206 M	$112 M
Insurance	5	$71 M	
Life Sciences	48	$160 M	$41 M
Manufacturing	678	$83 M	$20 M
O&G	100	$416 M	$54 M
Retail	49	$30 M	$12 M
Telecomm	420	$211 M	$30 M
Transportation	377	$1,804 M	$16 M
Grand Total	**2,335**	**$463 M**	**$44 M**

Source: Hyperion Research, 2019

New Business Markets in HPDA, AI and Deep Learning

The high performance data analysis (HPDA) market represents the convergence of long-standing, data-intensive modeling and simulation (M&S) methods in the HPC industry/application segments that Hyperion Research has tracked for more than 25 years, and newer high performance analytics methods that are increasingly employed in these segments as well as by commercial organizations that are adopting HPC for the first time. HPDA may employ either long-standing numerical M&S methods, newer methods such as large-scale graph analytics, semantic technologies, and knowledge discovery algorithms, or some combination of long-standing and newer methods.

Hyperion Research forecasts that HPDA server sales will reach $4.6 billion in 2022 (14.9% CAGR) to represent about 23% of the $19.6 billion figure for all HPC server revenue (Figure 3). Within the HPDA portion of the HPC realm, the market for HPC server-based artificial intelligence (AI), including machine learning and deep learning, is slated to grow at an even faster 26.3% CAGR to top $1.6 billion in 2022. HPC has already moved to the forefront of R&D for AI and deep learning, including leading-edge developments for autonomous vehicles, precision medicine, smart cities and the Internet of Things. These economically important new use cases are natural extensions of activities HPC has supported for some time—vehicle design, health care, urban transportation, regional/national power grids, and large distributed networks.

FIGURE 3

Worldwide HPC-based HPDA and AI Server Forecast ($ Millions)

	2017	2018	2019	2020	2021	2022	CAGR 17-22
Total WW HPDA Server Revenues	2,333	2,830	3,224	3,488	4,040	4,680	14.9%
Total HPC-Based AI (DL, ML, and Other)	501	673	845	986	1,260	1,611	26.3%

Source: Hyperion Research, 2019

CHALLENGES FOR HPC SYSTEMS HAVE GROWN

Sustained, Balanced Performance: The Holy Grail for HPC Users

For years, many HPC architectures have tilted further and further away from a reasonable balance between processor speed, memory access, and I/O speed. As successive generations of HPC systems have upped peak processor performance without corresponding advances in per-core memory capacity and speed, systems have become increasingly compute-centric, and the well-known "memory wall" has gotten worse.

Adding to this dilemma is the new HPC Big Data era with its demands for superb memory and I/O capabilities, sometimes with little need for computing prowess or double precision. Data-intensive (high-performance data analysis) workloads are emerging relatively rapidly. These emerging data-intensive problems are exposing more limitations of established HPC architectural designs, not just the memory wall itself, but also how existing compute-centric architectures handle data movement throughout the system. It's important to make advances here, or data movement for AI, deep learning and other

emerging high-performance data analysis problems could become frustratingly slow and expensive for users in the business, government and academic sectors.

The Hardware Challenges

Although businesses rarely need to acquire the world's largest supercomputers, it's instructive to see how fast the processor-driven peak performance of this class of supercomputers has grown–since the same trend has affected other HPC systems. As Table 3 shows, on the June 2000 list of the world's most powerful supercomputers (www.top500.org), the number one system boasted 9,632 cores and peak performance of 3.2 teraflops (TF). Fast forward to the June 2018 list and the number one supercomputer featured about 2.3 million processor cores and 187.7 peak petaflops (PF). That's a 228-fold increase in the core count and a 58,643-fold peak performance leap.

TABLE 3

Growth in System Size: Top Supercomputer on Top500 List

	June 2000	June 2018	Gain Factor (X)
Peak Performance (TF)	3.2	187,659	58,643
Core Count	9,632	228.3	236,974
Peak Performance/Core (GF)	0.33	82.2	246.6

Note: The data is from www.top500.org

Source: Hyperion Research, 2017

It's clear that to deliver the sustained, balanced performance that creates strong productivity and TCO, HPC systems need to focus on much more than raw processor speed. They also need to be memory-driven, providing adequate memory capacity to handle today's more data-intensive workloads and sufficient memory speed to keep all those processor cores continuously fed with data.

Packing a far larger average number of processors, memory parts and other components into today's HPC systems creates other major challenges. To fit into customers' existing data centers and avoid or delay the need for new data centers, HPC systems need to be much denser than ever before. But bigger, denser HPC systems generate more heat, giving rise to the trend toward liquid cooling as a more-efficient alternative to air-cooled systems.

The Critical Role of System Software

Although imbalance (extreme compute-centrism) is arguably the biggest performance problem facing HPC systems today, it is not the only major problem. Another important issue for buyers and users is managing the complex, growing set of requirements for operating these systems. These mounting requirements have made it harder to present users with an HPC resource that is comprehensive, coherent, tightly integrated, and highly performant on their applications.

The system management requirements fall into these main categories:

- Heterogeneous workloads (floating point-based simulation, integer-based analytics)

- The movement from synchronous applications to asynchronous workflows
- Rapid growth in average system sizes and component counts
- Heterogeneous processing elements (CPUs, coprocessors/accelerators)
- Heterogeneous environments (on premise data centers, public/private/hybrid clouds)
- Reliability/resiliency at scale
- Power efficiency and power awareness ("power steering")
- Cyber security

The chief responsibility for managing this complexity falls to the HPC software stack, i.e., the software between the hardware and the applications layer. During the past decade, HPC stacks have had to become much more sophisticated and intelligent. But the evolving list of requirements shown above presents major new challenges for stack developers and end users.

Storage: Another Major Challenge

HPC storage is increasingly playing a critical role in determining the overall performance and utility of a complete HPC system. There are a number of interrelated trends contributing to this:

- Data access is the primary limiting factor in HPC performance. Almost all processors today and going forward realize their increasing rates of computational performance by adding additional parallel computational units. As such, the demand for more data at higher transfer rates and with lower wait times for reads and writes is shifting the burden of overall system performance away from the processor to the memory/storage system.
- The case is complicated by the proliferation of HPC designs that rely heavily on specialized processors such as graphic processing units (GPUs) or other custom accelerators that perform a limited class of specialized, typically, data intensive operations.
- New HPC use cases are driving a rapid increase in the types of data used in HPC systems. Unlike the early days of HPC, where most of the data involved in a program consisted primarily of regular, standardized, and uniformly formatted numerical data sets, more and more HPC applications are relying on an increasingly diverse set of data, each of which must be handled efficiently and effectively by the data storage infrastructure.
- Data structures that are becoming more prevalent include structured and unstructured data, audio, video, medical and other precision instrument output, advanced manufacturing equipment output, and a data flows from myriad IoT devices.
- Increasingly, HPC jobs will to need to support a mix of different data types within a single job.
- Likewise, new HPC use cases are broadening the base of requirements to which a well-balanced storage system must respond. Of particular significance is the growing demand for real-time analysis, such as in financial or fraud detection applications, new and typically unpredictable data access patterns for critical big data jobs such as graph analytics, and emerging requirements for heterogeneous data such as those involved with analytics applications drawing on social media inputs.

In response to these trends, HPC designers increasingly look to implement new storage/memory schemes to speed performance. The inclusion of new features such as multi-level caches, burst buffer schemes, solid state disks, flash memory options, network supported data processing capabilities, and disk options that span the price/performance spectrum are all becoming an integral part of the HPC storage mix. In such architectures, the ability to control the flow of data across a single job, system, or even a data center is becoming increasingly challenging. Optimizing - or indeed simply managing – that

flow in a real-time environment will increasingly become the responsibly of a sophisticated hierarchical storage management system.

As the next section of this report explains, HPE has developed a strong portfolio of its own HPC systems that include purpose-built system management, an application acceleration suite, and storage data management solutions. At the same time, with partners HPE also offers a comprehensive "best of breed" portfolio of HPC management solutions, including workload management, software development tools and more.

EXAMINING HPE IN THE HPC MARKET

In light of the market opportunities and challenges discussed above, this section looks at HPE's historical rise in the global HPC market and the company's more recent success in the market for the most capable leadership class of supercomputers – represented by HPE Apollo and HPE SGI 8600 systems.

Hewlett Packard Enterprise Is the HPC Market Leader

Hewlett Packard Enterprise has been the biggest OEM driver and beneficiary of HPC market growth. In 2003, HPE (then part of HP) slightly outpaced IBM to become the global HPC market leader by posting $1.8 billion in HPC server revenue. Fast forward to full-year 2016 and HPE more than doubled its total to $3.9 billion, or 34.6% of the global market, not counting the contribution from HPE's acquisition of SGI. With SGI revenue added on a pro forma basis for both years, HPE's HPC server revenue rose from $2.0 billion in 2003 to $4.2 billion in 2017. That equaled 34.2% market share, compared with 19.0% for HPE's nearest rival. HPE also achieved strong performance in 2018, Antonio Neri's first full fiscal year since he took the helm from Meg Whitman. The company reported higher-than-expected sales growth of seven percent to $30.9 billion, including 25% year-year growth in high performance computing.

The HPE Apollo Product Lineup

HPE Apollo Systems are high-density, energy-efficient, memory-intensive systems that scale from midrange to leading-edge supercomputers. Standout characteristics include massive scale-out range and advanced water cooling (no exotic liquids), designed to allow data centers to fit more compute and memory power into their power and spatial envelopes.

Select HPE Apollo systems leverage the latest Intel Xeon Scalable processors. Intel's platform innovations are aimed at the convergence of HPC and AI, and are supported by a broad software ecosystem, along with lower latency and greater capacity for improved efficiency. These systems also exploit that Intel Omni-Path Architecture (Intel OPA) that's designed as a low latency interconnect for scalable performance in multi-node environments, such as AI training applications.

Following is the HPE Apollo compute lineup:

- **HPE Apollo 2000 Gen10 system** is positioned as the enterprise bridge to scale-out architecture and supports a variety of workloads, from remote site systems to large HPC clusters and everything in between. This density-optimized, 2U shared infrastructure chassis supports up to 4 ProLiant XL170r Gen10 and up to 2 ProLiant XL190r Gen10 independent, hot-plug servers with the latest Intel Xeon Scalable processors. It includes all the traditional data center attributes: standard racks and cabling and rear-aisle serviceability access. A 42U rack fits up to 20 Apollo r2000 series chassis accommodating up to 80 servers per rack.

- **HPE Apollo 6500 Gen10 system** is an HPC and deep learning platform with performance bolstered by GPUs, a fast GPU interconnect, high bandwidth fabric and a configurable GPU topology to match users' workloads. The system's reliability, availability, and serviceability (RAS) features include up to 8 GPUs per server, NVLink for fast GPU-to-GPU communication, the latest Intel Xeon Scalable processors support, choice of high-speed / low latency fabric, and is workload enhanced using flexible configuration capabilities. While aimed at deep learning workloads, the system is also designed to excel on complex simulation and modeling workloads.
- **HPE Apollo 6000 Gen10 system** is designed with density optimization and HPC solution integration to help manage and scale HPC computing demands. The modular, fully integrated, flexible HPE Apollo k6000 Chassis accommodates up to 24 ProLiant XL230k servers with Intel Xeon Scalable processors and redundant power and cooling, along with ethernet and high speed fabric switches to address varied HPC workload needs. This Apollo system is a performance optimized, air-cooled solution designed for customers with space limitations. This solution offers configuration flexibility, simplified administration and rack-scale efficiency, all aimed at lowering the total cost of ownership for mass scale HPC deployments.
- **HPE Apollo 70 server** is an Arm-based production-ready HPC system. The Marvell ThunderX2 processor offers up to 32 cores and strong memory bandwidth per core for optimal performance on memory-intensive HPC workloads. The HPE Apollo 70 is used in the world's first Arm-based supercomputer at Sandia National Laboratories, providing performance exceeding 1 Petaflop.

The HPE SGI 8600 System

HPE SGI 8600 system offers petaflop speed and scalability to thousands of nodes in a dense, proven architecture which currently powers some of the most powerful supercomputers in the TOP500. This supercomputer is a liquid cooled, tray-based, high-density clustered computer system designed from the ground up to deliver exceptional performance, scale, and density. The basic building block of the HPE SGI 8600 System is the E-cell. The E-cell consists of two 42U high E-racks which are separated by a cooling rack. The E-cell is a sealed unit, uses closed-loop cooling technology, and does not exhaust heated air into the data center. A direct attached liquid cooled "cold sink" provides for efficient heat removal from high power devices, including processors, GPUs, and switches via an auxiliary cooling distribution unit.

HPE Gen10 Servers: Strong on Security

HPE Gen10 servers offer comprehensive security features, down to the silicon. HPE Secure Compute Lifecycle, which includes the company's purpose-built iLO5 silicon chip, is designed to provide unexcelled protection, detection, and recovery capabilities. HPE engaged InfusionPoints to conduct an independent, comparative assessment of the security of the Gen10 Server line hardware and platform firmware against three of HPE's industry competitors. Tests included attacks against physical interfaces, platform firmware, and network interfaces. InfusionPoints reported that the HPE Gen10 server line is a significant step ahead of its competitors.

HPE Adaptive Rack Cooling System

Realizing that power requirements continue to increase, HPE offers the Adaptive Rack Cooling System (ARCS) as a solution to control the heat generated by HPC infrastructure. Not only do increasingly powerful servers generate additional heat, but storage and networking infrastructure will also need heat management. ARCS offers a closed loop liquid cooling system designed to cool up to 150KW of power while providing room neutral air. By using a quickly deployed warm water system, the HPE Adaptive

Rack Cooling System allows existing data centers to maximize their density using the latest generation infrastructure.

HPE's Software Solutions for HPC Users

Hyperion Research has said for years that software continues to be the number one roadblock to the adoption of HPC. Developing, installing, and running an HPC environment comes with many challenges, especially those listed earlier in this paper. More specifically, Hyperion Research has been pointing out that software, including systems management software will be even more important than hardware in driving future HPC progress, including advances in the HPDA-AI sectors.

HPE addresses these obstacles by offering a comprehensive, modular software portfolio combining HPE-engineered solutions with best-of-breed solutions from commercial partners and open source providers (such as leading workload management tools and software development tools). HPE validates, integrates and performance-optimizes all software solutions so they offer strong performance on HPE HPC systems.

HPE created the HPE HPC software portfolio to address the fundamental features users expect in a leading-edge software as well as important additional capabilities, including resiliency/security, operational efficiency, HPC-class scaling capabilities, and optimized HPC performance.

HPE Performance Cluster Manager

Managing HPC clusters is challenging, especially dealing with large-scale systems with thousands of nodes and other components to monitor and manage. System administrators need to manage complex HPC installations and keep system maintenance events to a minimum for optimal productivity. To address these requirements, HPE offers fully integrated system management software for all HPE HPC systems: HPE Performance Cluster Manager provides secure system setup, hardware monitoring and management, image management and software updates as well as power management for systems of any scale, up to 100,000 nodes.

HPE Message Passing Interface (MPI)

HPE MPI is a development environment designed to enable the development and optimization of high performance computing applications. It leverages optimized software libraries, runtime tools, and a scalable development environment to help customers tune and accelerate compute-intensive applications running on any HPE Linux-based cluster.

HPE's Storage Solutions for HPC

HPE offers a variety of storage solutions to meet the data management and storage requirements for HPC computing. HPC often involves extremely large datasets and presents challenges for data management and throughput as well as data archiving. New workloads, such as genomics and deep learning, introduce a variety of new file types and sizes as well.

Data Management Framework

HPE Data Management Framework (DMF) is designed to facilitate HPC data management. With an automated rules engine, extensible metadata attributes and file system management, DMF is designed to automate and streamline data workflows, simplify dataset creation and processing, and optimize usage of storage infrastructure. DMF aims to reduce the overall time for processing by providing data when it's needed, along with enabling archived data to be easily brought back for usage, allowing all past and present data managed by DMF to function as a virtual data layer.

Scalable Storage for Lustre

HPE Scalable Storage for Lustre solution is a custom, modular Lustre storage solution that can be tailored to specific workloads. This solution is tested for reliability and scalability, and is based on the HPE ProLiant DL360 and DL380 servers and the HPE D6020 Disk Enclosure. An HPE validated version of Community Lustre is provided, featuring the ZFS file system for data integrity, data compression, and file system snapshots. The Integrated Manager for Lustre is included and provides simplified system management and tuning. File system management is available through the HPE Data Management Framework, removing the requirement to install the Lustre Robinhood engine with its separate tracking database.

WekaIO Matrix

WekaIO Matrix is a high-performance, flash-optimized parallel file system that's designed to provide low-latency performance with next-generation workloads involving small file sizes as well as mixed file types. Designed as a radically simple storage solution with the performance of all-flash arrays as well as the scalability and economics of the cloud, Matrix combines the performance of NVMe-based flash storage with an S3 connection to object or cloud storage for economical capacity storage. Data is automatically tiered between the two layers as needed. An intuitive graphical user interface allows quick and easy management of data, with a command line interface as an alternative, and an API providing integration with management stacks.

Services for HPC deployments

HPE's IT services organization offers a range of services for HPC users:

- **Advisory services from HPE Pointnext** deliver advice, a library of field-proven blueprints, and IP-based methodologies.
- **Professional services from HPE Pointnext** provide design and delivery of users' solutions, from HPE Factory Express HPC Integration Service to installation and startup to training

Operational services from HPE Pointnext provide approaches such as HPE Flexible Capacity and Datacenter Care. HPE Flexible Capacity service offers a public cloud experience designed to match the benefits of on-premises IT. This pay-as-you-grow solution responds to changes in demand without the usual procurement process.

CASE STUDIES

The following section contains brief case studies that illustrate the benefits of the HPE Apollo and HPE SGI 8600 systems.

Pittsburgh Supercomputing Center Brings Advanced Supercomputing To Non-Programmers

High performance computing can be a powerful tool for driving the profound breakthroughs needed to cure disease, protect the environment, and deepen our understanding of our universe. But to maximize its impact, HPC must be accessible to researchers and scientists. So, the Pittsburgh Supercomputing Center leveraged HPE technology to build Bridges: a supercomputer that aims to be as easy to use as a laptop. The Pittsburgh Supercomputing Center provides researchers with access to supercomputers

for solving the most challenging problems in science and engineering, including treating complex diseases and protecting the environment.

Detecting cyberattacks. Curing breast cancer. Predicting severe storms. A common thread among these vexing challenges is their computational complexity. No ordinary computer can crunch calculations on a large enough scale to help researchers understand them. And the alternative: that class of high-powered systems known as supercomputers tends to be expensive and complex to build and manage. PSC's traditional user base used to include only engineers, chemists, and physicists, people who typically brought some specialized HPC skills. Now, PSC is also bringing its resources to researchers who are experts in their fields but might never have programmed a computer. To reach these new communities, PSC delivers HPC muscle through the types of software interfaces researchers are accustomed to using on their laptop PCs.

A new era of advanced supercomputing

Designed by PSC and built, racked, tested to PCS's specifications by HPE Pointnext, Bridges leverages advanced technologies from HPE, Intel, and NVIDIA to usher in a new era of supercomputing to tackle some of the world's most complex challenges. Approximately 400 researchers are currently using Bridges, 153 U.S. transplant centers optimize kidney donation exchange and metagenomes are being studied to make wheat and rice more resistant to drought, disease, and pests.

Mercedes-AMG Petronas Motorsport Races Ahead

HPE has joined the Mercedes-AMG Petronas Motorsport team to help pair its high performance race cars with advanced IT engineering. In a sport where split seconds make the difference, HPE technologies and HPE Pointnext help this F1™ team gain new technology advantages to power its drivers around the track.

For Mercedes-AMG Petronas Motorsport, winning requires a flawless combination of sophisticated auto engineering and shrewd on-the-track strategy. The right IT is critical at every stage of the race—starting well before race day with the HPC systems required to power state-of-the-art engineering, R&D, testing, and simulations. With HPE on the team, Mercedes-AMG Petronas Motorsport can support state-of-the-art machine learning. The team can tackle data challenges more quickly and efficiently and gain better insight into its data. And this in turn delivers the reduced development time and iterations that lead to better team performance.

NCAR Boosts Atmospheric Research Three-Fold With HPE SGI 8600 System Powered By The Intel Xeon Processor

Weather and climate research place nearly insatiable demands on HPC systems. Think here of high-resolution ensemble models that that can incorporate tens of petabytes of data and added variables such as the carbon cycle and physics. Think of the need for brute-force computing speed to improve models used to forecast severe storms and other weather events before they occur, or to simulate 30,000 years of climate change in reasonable time frames. Because climate and weather simulations can be very long-running problems, think of the need for the HPC system to exhibit extreme reliability as well.

How much processing power does it take to model the weather on a planet? In order to understand the magnitude of such a task, researchers need to model the planet itself as well as all its terrestrial and related geospace systems. It's no small task, but that's what the National Center for Atmospheric Research (NCAR) is doing.

In order to handle that deep scale of research, the NCAR team knew it would need a system that could analyze and perform cross comparisons on approximately 30 petabytes of data coming from all parts of the globe.

When all the bids were reviewed, the team chose the HPE SGI 8600 System running on Intel Xeon processors, which can scale to more than 10,000 nodes.

With the right mix of HPE high performance computing software, NCAR can make the best use of their system. HPE Performance Software enables the team to streamline software maintenance and task automation with bare metal provisioning, version control image management, power management, and system health management. HPE Performance Software. Message Passing Interface (MPI) offers optimized software libraries, a high performance MPI environment and runtime tools. HPE built and tested the supercomputer off-site prior to installing it. The customer reports that system can do three times as much work with a three-fold shrinkage in floor space, compared to the prior system. NCAR also gets three times as much science done for the same energy footprint.

With the computational power now available to NCAR, the research team will be looking into weather beyond the Earth's atmosphere. "We have people focusing now on the impact of space weather on our planet and its impact on weather and communications, and well, everything," explains Anke Kamrath, director of the Computing and Information Systems Laboratory at NCAR. "Now we can look at things that are happening on the sun and try to find the predictability of those events. When is it going to happen next? What can we do to limit their impacts when we do have these predictions? That's where the science is."

FUTURE OUTLOOK

Hyperion Research forecasts that the global market for HPC server systems will continue growing at a robust CAGR rate to reach $19.6 billion in 2022 and that the overall HPC market that also includes software, storage and service will expand to $38 billion in that year. We believe the rapid rise of data-intensive simulation and analytics, including economically important new use cases for HPDA and AI/deep learning, will ramp up demand for leading-edge HPC systems that balance strong processors/accelerators with strong memory capacities and data transfer rates, along with advanced software stacks that enable the systems to operate efficiently and resiliently even at ultra-large scale (e.g., exascale). In our view, systems of this kind are future-facing, because they are the first to tackle emerging workloads and use cases that promise to become economically very important within five years, such as autonomous vehicles, precision medicine, smart cities and the Internet of Things.

HPE, already the revenue leader in the global HPC market, has positioned itself well to become a strong player in the market for leading-edge HPC systems with the company's HPE Apollo product series and HPE SGI 8600 system. As the brief case studies in this paper attest, HPE is already building a successful track record in this strategically important market. HPE's acquisition of SGI boosts the combined company's ability to pursue opportunities for the largest supercomputers, including exascale systems. HPE's CEO has expressed a long-term commitment to the HPC market and the company is engaged in serious R&D initiatives on multiple fronts. Last year, HPE was awarded a research grant from the U.S.

Department of Energy (DOE) to develop a reference design for a memory-driven exascale supercomputer.

HPE is an active member of the Gen-Z consortium that aims to develop a new universal interconnect designed to provide memory semantic access to data and devices via direct-attached, switched or fabric topologies that will enable simpler and more powerful computer architectures. Gen-Z's efficient, memory-semantic protocol simplifies hardware and software designs, and supports a wide-range of signaling rates and link widths that enable solutions to scale from 10s to 100s of GB/s of bandwidth. Another important, sometimes under-recognized contributor to HPE's success is Hewlett Packard Labs, which conducts world-class research on topics highly relevant to the HPC community, such as characterization and benchmarking of deep learning problems and solutions.

In summary, Hyperion Research believes that the HPE Apollo series and the HPE SGI 8600 system boost the HPC market leader's potential for additional growth in the expanding global market for high performance computing systems.

About Hyperion Research, LLC

Hyperion Research, consisting of the former IDC high performance computing (HPC) analyst team, provides HPC information, analysis, and recommendations based on technology and market trends. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA/AI (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.hpcuserforum.com and www.HyperionResearch.com
