



**Hewlett Packard
Enterprise**

Brochure



The NSCI and Hewlett Packard Enterprise contribution



HPE ProLiant Gen9 servers
Powered by Intel Xeon processors.
Intel Inside®. Powerful Solution Outside.



Executive summary

The President of the United States issued an executive order establishing the National Strategic Computing Initiative (NSCI) on 29 July 2015. The order outlined two primary rationales for establishing the NSCI: retaining the United States' leadership in high-end computing in the face of increasing competition from foreign interests, and maximizing the impact of that leadership through increased competitiveness by the U.S. scientific and industrial outcomes through widespread adoption of computational approaches.

Hewlett Packard **announced a corporate separation in October 2015** that defined two new businesses, one focused on providing personal computing products and services (HP Inc.), and the other, Hewlett Packard Enterprise, providing leading infrastructure, software, and services to enterprise and public sector customers. As part of their overall product portfolio, Hewlett Packard Enterprise has a concentration on High Performance Computing (HPC), allowing it to contribute broadly to the NSCI.

The NSCI as a broad-based initiative

The executive order gave specifications and timelines for each of the lead agencies: Department of Defense (DOD), the National Science Foundation (NSF), and the Department of Energy (DOE). It also delineated areas of responsibility amongst those agencies. In brief, the DOD will work on analytics, the NSF on “beyond Moore’s Law” compute technology, workforce development, and greater proliferation of HPC throughout the public and private sectors, and the DOE on exascale-level computing for modeling and simulation. The agencies have 90 days from the issuance of the order to respond with a plan on their collective approach to the NSCI’s objectives for the next 15 years.

A strategic commitment to HPC

As a part of its separation, Hewlett Packard Enterprise designated HPC and Big Data as targeted growth opportunities and subsequently created a specific Global Business Unit (GBU) to pursue those market segments. This has focused investment and created end-to-end responsibility for strategy, product development, and services tailored specifically to meet the needs of customers in this space. The GBU hosts an engineering team that will focus on designing and building HPC and Big Data optimized solutions. The GBU will expand its domain expertise and application-performance skills necessary to effectively engage with customers to solve their most pressing challenges.



The GBU has the full support of the CEO and board of directors and is both chartered and resourced for growth. Hewlett Packard Enterprise has tremendous scale across the full spectrum of capability necessary to contribute to the NSCI with one of the largest, most successful IT research institutions that is well connected to the company's development capability in servers, storage, and software. HPE has an unmatched supply chain and factory base that can build and deliver complex systems through its own world-class HPE Services organization. HPE has a direct salesforce that covers the full range of vertical industries with trusted advisor relationships working in close partnership with the GBU for both development and delivery of customer solutions.

Ramping up current contributions to HPC

Hewlett Packard Enterprise contributes significantly to HPC today and remains in a leadership position, with roughly a third of the total HPC market, based on IDC's market analysis. HPE's leadership position has grown even faster as a result of market restructuring. A good example is the divestiture of IBM's X-Series business.

Hewlett Packard Enterprise has significant market-segment strength in the manufacturing, energy, and financial services markets. Every major car company, aircraft manufacturer, and jet engine producer uses HPE servers in their design processes. In the energy segment, our solutions power discovery in the vast majority of oil and gas producers. Leading financial institutions in New York and Chicago run on HPE hardware. In the electronic design automation (EDA) domain, the chips that power the majority of HPC systems were designed on HPE systems. Indeed, our High Performance Computing solutions fuel many of the commercial economies of the world as shown by our longstanding leadership in the top-500 list of the world's most powerful computers.

Hewlett Packard Enterprise has deep expertise in building HPC clusters that match specific needs across many application segments. Teams of engineers and scientists provide application-level expertise, guiding customers to the most efficient configurations, maximizing performance and accessibility of HPC assets.

Hewlett Packard Enterprise also has a team of developers working on an HPC runtime software stack, scheduled to debut in 2016. The first phase, using open-source technology and guided by public sector customer engagement, will debut in 2016. This software stack will complement the industry-leading HPE Insight Cluster Management Utility, which provisions, controls, and monitors clusters of all sizes with unmatched functionality and performance. Follow-on software stacks will have additional industry-specific functionality and will sample from the best of available technologies. These will deploy over the years closing the decade.

Hewlett Packard Enterprise has a strong track record of identifying promising new technologies early in their innovation cycle and investing to bring them to maturity. One example is the high-performance parallel file systems where HPE was the prime contractor for Lustre Development through the U.S. Department of Energy. Accelerators increasingly power the HPC clusters at many institutions, whether they are GPUs, FPGAs, or multi-core coprocessors. HPE engaged early with leading suppliers and has a team of engineers focused on optimizing the use of accelerators, developing the broadest range of hardware options for deploying accelerators of any vendor. HPE is recognized as a leading early adopter of system-on-chip technology (examples in HPE Moonshot products) and is at the forefront of both non-volatile memory (**Persistent Memory**) and **silicon photonics**.





Broad market approach provides interesting technology to HPC

Hewlett Packard Enterprise is involved in the broader server market and is the leading server provider in the world today. Many of the latest technological changes are now occurring in these other markets that can have a home within the HPC landscape. HPE leads in adjacent markets also, such as the ISP segment, which continues to grow rapidly and demands quick technology development and adoption. Modular computing centers (containers that allow for rapid deployment of IT assets) first debuted at-scale in these markets and are now also deployed in HPC environments around the world. These customers are also at the forefront of highly scalable object storage, driving rapidly evolving solutions that can be brought to bear on the most data-demanding HPC use cases.

Scaling accessibility to HPC and Big Data solutions, solving the “missing middle” problem, and more

The “missing middle” phenomenon was first identified during the outset of the Blue Collar Computing initiative at the Ohio Supercomputing Center (OSC) in 2003. The observation was that small- and medium-sized businesses (SMBs) as well as most small to mid-sized academic institutions were not availing themselves of HPC even though it would be of great benefit to do so. Serious questions were raised regarding the cause of this phenomenon. Through a closely coordinated effort with Intel® and OSC, Hewlett Packard Enterprise helped identify scalable solutions to the most challenging HPC-adoption obstacles. This work resulted in language being added to the first renewal of the America COMPETES Act, which intended to initiate a national agenda to drive it top-down. Though much has been learned, the strategy is yet to be applied at scale. HPE, with its OSC and Intel partners, is best positioned to continue these efforts in earnest.

Cloud computing is an increasingly common method for deploying IT resources by many users. Hewlett Packard Enterprise has developed or acquired many cloud assets, including HPE Helion OpenStack® cloud software and is a leading contributor to the OpenStack initiative, particularly in the area of object storage.



The broad base of HPE solutions allows potential users to utilize their HPC in whatever way they wish as opposed to being forced into an on-premises design. For example, HPE Enterprise Services provides CaaS services for a variety of customers in the manufacturing and banking industries. This allows their full HPC resources to be provided “on tap” or in a burst capability for high-demand time frames. HPE Flexible Capacity, a solution provided by HPE Technology Services, deploys HPE-owned and managed capacity within a customer’s data center and allows short- or long-term HPC capability without the risk to security and performance or resorting to a public cloud.

The HPE Apollo server line: purpose-built for HPC and Big Data

In the spring of 2014, HP (now Hewlett Packard Enterprise) announced the Apollo server family, specifically designed for HPC and Big Data, with use-case optimized designs segmented into the HPE Apollo 2000, 4000, 6000, and 8000 server lines.

Recognizing that “one size does not fit all,” these individual-yet-complementary solutions address distinctly different optimization choices for power, performance, and function. The HPE Apollo 2000 enhances the accessibility of HPC by optimizing performance for the budget while also focusing on ease of deployment. The server has flexible options that tailor to specific workloads, including those that call for traditional computing mixed with accelerated computing. The HPE Apollo 6000 includes the functionality of the Apollo 2000 with additional optimizations for power efficiency and system reliability, which become necessary for large-scale systems. The HPE Apollo 8000 adds optimizations for density, bringing as much as three times the computational capacity per square foot as traditional solutions, and further adds optimizations for power efficiency, enabling PUEs well below 1.1. The HPE Apollo 4000 complements all of the other lines with optimization for the high-capacity and high-performance storage required in modern HPC and Big Data solutions.

Each of these server lines is software-compatible to the others, and they are frequently combined in deployments, giving maximum flexibility to develop a complete solution with optimizations wherever they are needed. All of the HPE Apollo servers were designed by a team with combined experience in developing dense, highly available designs and developed in the U.S. at HPE’s Houston-based engineering facility.

The intersection of HPC and Big Data

HPC is typically used to process data so it can then be further analyzed, stored, or forwarded to other devices for use or additional analysis. HPC has been used to manage Big Data problems for years. At the same time, the power of HPC resources allows many permutations of a particular simulation to be made using different data sets or boundary conditions. Each of these simulations can then be further analyzed using Big Data techniques to find patterns within them.

The strength of our solution strategies has enabled HPE's HPC and Big Data solutions to successfully meet the highly-scaled and business-critical needs of our customers. HPE solutions power the world's largest retailer. Even with these accomplishments, we remain resolved in our search for innovation to meet the specific needs of additional use cases as well as the needs of our existing customers as they evolve.

The HPE-Intel HPC Alliance

Hewlett Packard Enterprise and Intel Corporation have a long history of partnership in technology. As Intel is the world's leading processor supplier and HPE the leading HPC provider, it only makes sense that the two should work together in the HPC market. In July 2015, HP (now Hewlett Packard Enterprise) and Intel announced an HPC Alliance to bring together a broad range of technologies with a solution-oriented, vertical focus as opposed to purely a product or technology focus. The HPC Alliance includes the Intel Scalable System Framework of technologies for HPC including OmniPath fabrics and Intel® Xeon® Phi coprocessors. Moving forward, HPE and Intel will jointly collaborate in creating best-in-class HPC solutions. Engineering teams from HPE and Intel meet weekly to discuss future designs and get the best of each other's technologies for the benefit of HPC users.

The HPC Alliance will create solutions by working closely with industry-leading customers in government, financial services, energy, and life sciences. The resulting solutions will then be leveraged more broadly to allow a best-in-class solution for particular workloads. HPE and Intel fully expect this partnership to accelerate the adoption of new technology in concert with the goals of NSCI.

The capstone of the HPC Alliance is the creation of a Center of Excellence (CoE) based in Houston. The CoE allows HPE and Intel to work together on their latest technology while making it available to customers to explore, sample, and get the best out of their workloads and data sets as early as possible. A customer may bring their codes into the CoE (physically or virtually) and modernize and tune for the latest technologies from HPE and Intel.

The High Performance Consortium for Advanced Scientific and Technical Computing

Hewlett Packard Enterprise also works closely with customers in academia, government, and enterprise to further the effectiveness of HPC. The High Performance Consortium for Advanced Scientific and Technical Computing user group (HP-CAST) works to increase the capabilities of HPE solutions for large-scale, scientific, and technical computing. HP-CAST provides guidance to HPE on the essential development and support issues for such systems. HP-CAST meetings typically include corporate briefings and presentations by HPE executives and technical staff and discussions of customer issues related to high-performance technical computing.

Hewlett Packard Labs future contribution to HPC

Hewlett Packard Enterprise's broad appeal will accomplish many of the objectives of the NSCI, while the company's technology capability will future-proof many solutions going forward. Well before the issuance of the NSCI, Hewlett Packard Labs began researching how to grow computational capacity and accessibility while greatly enhancing performance per unit energy through rethinking how a computer is designed.

In this work, Hewlett Packard Labs engages in groundbreaking innovation around the architecture of distinct memory and computational units. This has led to leadership research in non-volatile random access memory (NVRAM, based on Persistent Memory technology), and photonics embedded in silicon devices (silicon photonics). Hewlett Packard Enterprise is pursuing a vision called "The Machine," which combines these technologies with processing elements and extremely fine, granular, data access capabilities. In The Machine, enormous memory pools can be created with extreme connectivity to have data become closer to everything. With this design, data motion can be both rapid and efficient. This unification of memory with communications fabric creates a new computational paradigm.

HPE recognizes that the pursuit of this vision also must incorporate the software redesign necessary to harness the new efficiencies that The Machine offers. HPE research also includes pathfinding of the systems software and application environments for The Machine. This pathfinding is done in partnership with key customers who can offer insight into their future computational needs and code to help migrate to The Machine. To support this work, prototypes will be deployed in 2016.

The Machine concepts are also aimed at machine learning and broad-scale Big Data use cases. With the unified memory and fabric paradigm, the development of data-to-thought perception and intelligent data mining will become much easier to accomplish. The Machine activity is well timed with the objectives of the NSCI mission objectives.

HPE as a full-service provider for NSCI

In addition to strengths in solutions and research, Hewlett Packard Enterprise provides a wealth of additional capabilities relevant to the NSCI mission, including:

- U.S.-based manufacturing and integration capability, based in Houston, Texas, which allows trusted production and deployment of systems for secure government use
- Enterprise and Technology Services allows users to explore and sample HPC in the format that best works for them and their budgets, whether on-premises or otherwise
- Myriad additional technologies being developed for many different workloads at many different scales that can be applied to HPC in the future such as virtualization and remote visualization
- Significant software development resources for HPC, Big Data, and cloud applications
- A very large supply chain allowing for best-in-class capabilities and pricing across a broad range of technologies
- The financial soundness and security delivered by a large profitable corporation with revenues of approximately \$50 billion



NSCI and HPE

The NSCI provides a way forward for the United States to invest in key technologies that will allow it to secure its borders, its interests overseas, and its competitiveness in an increasingly competitive world.

HPE is a U.S.-based company providing substantial technological intellectual property as well as jobs for tens of thousands of Americans. Hewlett Packard Enterprise will provide a unique set of capabilities to deliver on the promise of NSCI and has a specific focus around HPC and Big Data that will allow it to continue delivering cutting-edge technologies for a wide variety of users.

Learn more at
hpe.com/info/apollo



Sign up for updates

★ Rate this document



© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel, Intel Xeon, Intel Inside, and Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries.

4AA6-4718ENW, April 2016