Hewlett Packard Enterprise

Case Study

PITTSBURGH SUPERCOMPUTING CENTER

PSC revolutionizes supercomputing with HPE and Intel

Innovative HPC solution advances scientific research on a massive scale

Objective

Support diverse communities of researchers involved in both traditional and non-traditional high-performance computing projects

Approach

Create a whole new approach to supercomputing by combining high-performance computing (HPC) and big data in a fully integrated computing resource, utilizing very large-memory compute nodes and large numbers of traditional x86 as well as GPU-accelerated compute nodes, all interconnected across a high-speed fabric

IT Matters

- Increases aggregate shared memory available for big data research by over 8.6 times compared to previous HPC platform
- Delivers more than 36 times greater compute capacity to accelerate projects compared to previous HPC platform
- Achieves 50 times greater peak I/O bandwidth to parallel storage compared to previous HPC platform

Business Matters

- Enables ground-breaking research for understanding and treating complex diseases
- Runs advanced computer models that mimic human abilities, accelerating innovations such as robotics and speech recognition
- Extends supercomputing to new research communities, enabling them to scale their research to levels never before possible

Pioneering the world of supercomputing

Behind the curtain of everyday life there exists a world populated with people focused not on "what is" but "what is possible." They are constantly pushing the edge of human understanding to discover new ways of improving the human condition. They are researchers—explorers in fields as diverse as neuroscience, economics, social sciences, biology, and computer science.

At the core of their advanced research is the monumental task of problem solving. It requires running extraordinarily complex algorithms to figure out things like how the brain works, what causes disease, or where the next severe storm will strike. No ordinary computer is capable of handling the necessary calculations.

That's why researchers from across the U.S. turn to the Pittsburgh Supercomputing Center (PSC), a joint effort of Carnegie Mellon University and University of Pittsburgh. For more than 30 years, PSC has served the needs of the research community with cutting-edge computing capabilities. Always looking for innovative ways to add value for researchers, PSC first deployed a Cray supercomputer in 1986; it deployed the first massive parallel processing (MPP) systems; and in recent years, PSC has been home to the world's largest shared-memory system.

As a supercomputing pioneer, PSC has recently taken another giant leap. By combining high-performance computing (HPC) and big data in a fully integrated computing resource, PSC is changing the game once again for researchers. The new system, called Bridges, is built on large-memory and high-performance servers from

"We chose to work with HPE because they could provide the best set of components for what we hoped to achieve. We brainstormed extensively with HPE to optimize our combination of large-memory Superdome X and ProLiant servers with their compute-intensive Apollo systems. Bridges embodies unprecedented heterogeneity, and HPE's unique depth ultimately made it possible."

— Nick Nystrom, Senior Director of Research, Pittsburgh Supercomputing Center

Hewlett Packard Enterprise (HPE), deployed with the first-ever production implementation of Intel's Omni-Path Architecture (OPA).

One of the hallmarks of Bridges is its flexibility to support both traditional HPC users, such as engineers, chemists, and physicists, as well as new communities of researchers who never had access to HPC resources before. And that's opening huge opportunities in areas such as genome sequencing, machine learning, and even economics.

Nick Nystrom, senior director of research at PSC, remarks, "PSC has always embraced being the first to deliver new capabilities to research, and Bridges is no exception. Unlike traditional supercomputing, Bridges is designed to serve many different research communities with a heterogeneous architecture optimized for performance, flexibility, and ease-of-use. It's empowering research to enable new data-driven insights in ways never before possible."

The next generation of HPC

With insights gained from its close interaction with the national community, PSC recognized that many types of research were hindered due to the lack of memory and system balance. Think of the social scientist searching for a breakthrough in economic indicators with only 8 or 16GB of RAM on her laptop. Or imagine having the world's most powerful microscopes capturing vivid new images of the brain, but no way to process the terabytes of data collected every hour.

With a $17.2 million grant from the National Science Foundation (NSF), PSC set out to blow the lid off "what is possible." Its team of computer architects and computational scientists quickly drew up the blueprint for what would become Bridges. The question was whether the current state of information technology could support the concept.

"We knew what we needed," says Nystrom, "and we're familiar with what's available out in the marketplace. We chose to work with HPE because they could provide the best set of components for what we hoped to achieve. And they came to us as a real partner, not just a vendor. PSC pushed hard, and we brainstormed extensively with HPE to optimize our combination of large-memory Superdome X and ProLiant servers with their compute-intensive Apollo systems. Bridges embodies unprecedented heterogeneity, and HPE's unique depth ultimately made it possible."

PSC engaged HPE Technology Services to provide Installation and Deployment Services for the HPE infrastructure. This includes 4 HPE Integrity Superdome X servers with 16 Intel Xeon EX-series CPUs and 12TB of RAM each, 42 HPE ProLiant DL580 Gen9 Servers with 4 Intel Xeon EX-series CPUs and 3TB of RAM each, and 800 Apollo 2000 nodes with two Intel Xeon EP-series CPUs and 128GB of RAM each. Of the Apollo 2000 nodes, 48 also include NVIDIA Tesla GPU accelerators for computationally intensive workloads.

"I genuinely enjoyed working with HPE Technology Services team," observes Nystrom. "What we designed for integrating modern data analytics with high-end computing differed significantly from what one would build for traditional workloads. We talked about it with the HPE Technology Services team, they understood the value proposition, and they helped us configure a system that has proven to be a great success."

Stitching together the infrastructure is a 100Gbps fabric built on the Intel OPA. Six Intel OPA edge switches form the "core" of Bridges' custom topology, interconnecting the Superdome X and ProLiant DL580 servers with storage, database, web, data transfer, login, and management nodes. The core switches also link to 20 "leaf" switches, also implemented as Intel OPA edge switches, which connect the central topology to the Apollo 2000 nodes. With this groundbreaking infrastructure, any compute node has multiple paths to data in PSC's custom-built 10PB parallel file system, enabling high-speed workloads to run on a massive scale.

Enables ground-breaking new research

With HPE and Intel, Bridges is taking supercomputing to a whole new level. Compared to PSC's previous system called Blacklight, which was based on SGI UltraViolet computers and represented the world's largest shared memory system in its day, Bridges delivers over 8.6 times the aggregate shared memory and more than 36 times the compute capacity.

The beauty of Bridges is that it allows access to its diverse compute resources through a uniquely flexible software environment, allowing researchers who have never used high-performance computing to continue using paradigms that they use on their laptops. This includes using software such as Jupyter, R, and MATLAB, and it also includes "gateways," or domain-specific browser-based interfaces, that are already used by tens of thousands of researchers. This flexibility lets them work in the applications, languages, and research paradigms they're accustomed to, yet amp up their projects to supercomputing levels.

"It's Supercomputing-as-a-Service," quips Nystrom. "In many cases, our users only know that they're using a supercomputer by being able to do larger problems faster. We've designed Bridges so they can easily extend their current research into the supercomputing realm without the need to become programmers or learn a whole new environment. That lets them focus on what's most important—making discoveries."

One of the many research communities already taking advantage of Bridges is involved in genome sequence assembly. Researchers at the University of Georgia assembled the genome sequence of the human gut microbiome, which is a community of microorganisms that live inside all of us. To understand the enormity of this project, consider that the entire human genome is made up of 3 billion base pairs of DNA molecules. The human gut microbiome contains 378 billion base pairs.

The large-memory servers in Bridges are ideal for genome sequence assembly. These types of large-scale projects do not run effectively across multiple compute nodes. But a compute node with 3 or 12TB of RAM allows them to run without constraint. For the University of Georgia project, researchers

Case study
Pittsburgh Supercomputing Center

Industry
Scientific research

Customer at a glance

Application
- Research in digital humanities, social sciences, artificial intelligence, physical sciences, biology, neuroscience, business, public health, and computer science

Hardware
- HPE Integrity Superdome X Servers
- HPE ProLiant DL580 Gen9 Servers
- HPE Apollo 2000 System
- Intel Omni-Path Architecture

Services
- HPE Technology Services Installation and Deployment Services
- HPE Foundation Care

are now able to compare the microbiomes of normal healthy human guts with those in people with Type 2 diabetes. This is revealing how different microorganisms affect, and are affected by, diabetes, with the hope of eventually finding a cure.

Another group of researchers leverage the GPU-accelerated Apollo 2000 nodes for "deep learning" projects. These involve training neural networks—a computer system modeled on the human brain and nervous system—to recognize things the same way humans do. These types of computationally intensive problems require lots of very fast compute nodes. By running their calculations in Bridges using the Apollo 2000 servers with NVIDIA GPUs, researchers can bring the training time for their neural networks down from weeks to days, or in some cases days to hours. This allows new innovations to come into the market faster and faster. Today it's neuroscience and electronic markets. Tomorrow, it could be fully functional humanoid robots.

PSC also runs its own research projects on Bridges. For example, a group of researchers from the University of Pittsburgh, Carnegie Mellon University, and PSC are focused on identifying cause-and-effect relationships in three biological areas: cancer, chronic lung disease, and brain disorders. These involve very large data sets, including genomic, fMRI, and image data. In the past, the algorithms used to solve these cause-and-effect problems could only scale to about 100 variables. But the data sets contain up to hundreds of thousands of variables. Now with Bridges, thanks to the large memory afforded by Superdome X and ProLiant DL580 servers, the research team has been able to scale its algorithms to run one million variables, and is making important progress on these biological challenges.

Sustaining strong research progress and providing a good user experience require high reliability and availability. PSC reports that Bridges' components – both HPE servers and the Intel OPA interconnect – have been extremely reliable, and Bridges' architecture is such that if any component is offline, it affects only a very small part of the overall system. But with thousands of components, service is inevitably required, and for that PSC has HPE Foundation Care to supply rapid repair or replacement. The HPE Foundation Care professionals were there for Bridges' installation, know the system, and continue to provide first-rate service, working behind the scenes to ensure that research continues interrupted.

"Without Bridges and its HPE infrastructure, this kind of research simply could not happen," notes Nystrom. "The discoveries and insights our researchers are now uncovering will have direct effects on human lives by way of advancing precision medicine, increasing energy efficiency, and improving policy-making for the economy. And we've only just begun to explore all the capabilities Bridges has to offer."

Learn more at
hpe.com

Sign up for updates

a00002585ENW, February 2017,