Hewlett Packard Enterprise

Case Study

GITDA deploys supercomputers to support research throughout Hungary

Implements HPE supercomputers to meet HPC requirements nationwide

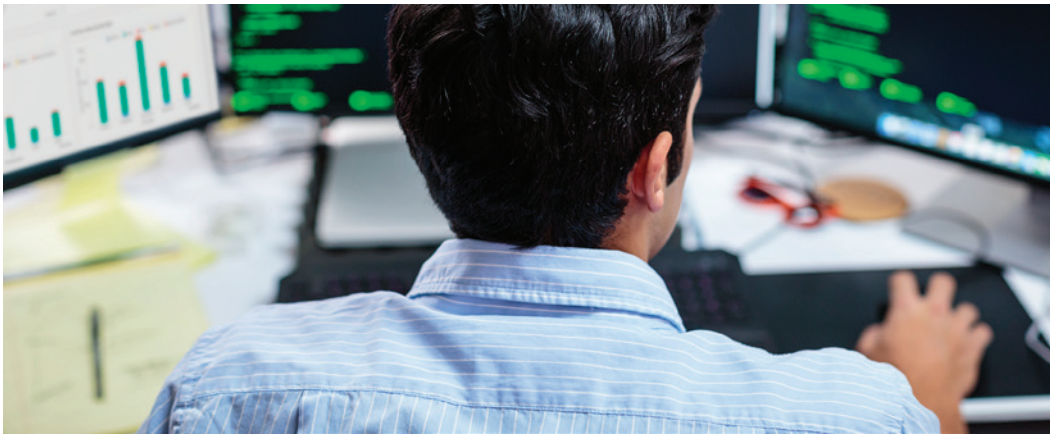

Objective
Provide heterogeneous supercomputing resources to enable the advanced computational demands of universities and research centers

Approach
Build and manage powerful and energy-efficient computing infrastructure based on proven, high-performance servers from HPE

IT Matters
- GITDA supercomputers support over 200 concurrent research projects
- Centralized supercomputing resources are made available to 14 universities and research centers
- Zero downtime since implementation

Business Matters
- Swiftly deploy supercomputing infrastructure to meet aggressive deadlines set by the European Union
- Enable qualified institutions to access supercomputing resources to help retain Hungary's world-class research capabilities
- Reduce costs while improving performance and increasing HPC capacity

The costs of deploying and managing supercomputer infrastructure is too high for most universities and research facilities, so many countries are deploying centralized supercomputers to foster high-performance research capabilities nationwide. The KIFÜ Government Information Technology Development Agency (GITDA) has developed two supercomputing architectures based on Hewlett Packard Enterprise (HPE) platforms to enable world-class research throughout Hungary.

Supercomputers provide the massive compute power that allows leading research institutions to run the simulations and analytics that are behind incredible breakthroughs in science and technology. Time is of the essence when trying to find a cure, predict the weather, or create the next game-changing innovation.

High-performance computing is essential for enabling world-class research, and delivering supercomputer infrastructure allows researchers and universities to swiftly analyse and model massive amounts of data. Supercomputers are valuable tools for academic and scientific research, but the cost of deploying and managing high-performance server resources can be prohibitive.

That's why countries like Hungary are building centralized supercomputing resources and sharing them with advanced universities and research facilities. Two of the largest high-performance computing (HPC) systems in Hungary are now providing supercomputing resources to universities and research organizations nationwide.

"We provide a flexible, high-performance infrastructure to fulfill our primary objective of offering efficient and compact computing facilities for our research community. The ability to quickly and reliably process very large volumes of data enables advanced and innovative research throughout Hungary."

— Zoltan Kiss, Head of GITDA's HPC Technology Department

This supercomputer program was initially developed by Hungary's National Information Development Institute (NIIF), which is now part of the KIFÜ Government Information Technology Development Agency (GITDA). According to Zoltan Kiss, Head of GITDA's HPC Technology Department, "Our users are based all over the country, and they have a diverse set of needs. We needed to deploy two HPC implementations connected over a common InfiniBand network so we could efficiently allocate resources to best meet the needs of our users."

GITDA sought to deploy two accelerated HPC capabilities, and received funding from the European Union to help support the deployment. But to comply with EU requirements, the facilities had to be deployed quickly. "We needed a cost-effective HPC solution that would meet the needs of most of our users, but we also needed a dense, energy-efficient solution that would best leverage the energy management capabilities of our data center," explains Kiss.

Deploying a Supercomputing Server Cluster

For the first implementation, GITDA selected a distributed server cluster based on HPE ProLiant SL250 Gen8 Servers and shortly thereafter deployed the HPE Apollo 8000 System. "Our HPC infrastructure is heterogeneous, and it allows us to carefully allocate supercomputing resources based on the needs of our users," states Kiss. "Both resources are highly accelerated and we have developed the capability to flexibly respond to the computational needs of researchers across Hungary."

According to Kiss, "We had a good history with HPE and have been using HPE supercomputers since 2010. We needed a cost-effective and easy-to-manage HPC solution that we could deploy quickly, and HPE offered proven HPC solutions with excellent performance-to-cost ratios."

GITDA first deployed an 84-node ProLiant SL250 Gen8 Server cluster accelerated by Nvidia graphical processing units (GPUs). Each node has Intel® Xeon® processors, 128 GB of RAM, and three Nvidia GPUs for accelerating server performance. The server cluster was implemented on schedule, and now provides the majority of the HPC processing capacity for GITDA users.

"When we completed the installation, this was the largest HPC deployment in all of Hungary," Kiss proudly states. "We were delivering powerful processing to users nationwide, and then within a year our HPE Apollo 8000 System was running live."

Dense and energy-efficient supercomputing infrastructure

The next phase was to deploy a more energy-efficient supercomputing solution, and after careful search management selected the HPE Apollo 8000 System. "We wanted to drive down our operational costs, reduce our rack space, and make our data center more energy efficient," explains Kiss. "Apollo also offers the ability for us to leverage warm-water cooling, which will result in even greater efficiencies for us in the future."

The HPE Apollo high-density server family is built for the highest levels of performance and efficiency. They offer rack-scale compute, storage, networking, power and cooling for HPC workloads. The HPE Apollo 8000 System leads the industry in teraflops per rack. This performance is the result of highly efficient liquid cooling that's delivered risk-free, supporting data center energy recycling and lowering energy costs.

The Apollo 8000 is rack-configured, with 45 nodes offering GITDA a 106-teraflops-per-second computing capacity. The system has 1,056 cores, and each node has 128 GB of RAM. The Apollo 8000 has 90 Intel Xeon Phi™ processors installed, allowing GITDA to extract extreme performance from highly parallel applications.

Intel Xeon Phi coprocessors are based on Intel Many Integrated Core (MIC) Architecture and enable dramatic performance gains for many of GITDA's most demanding research projects. "The HPE Apollo 8000 configured with Intel Phi processors is allowing us to achieve optimized performance for even our most highly parallel technical computing workloads while efficiently maintaining a unified hardware and software environment," explains Kiss.

Efficient management of supercomputing resources

Both supercomputer deployments share HPE 3PAR StoreServ 7200 and Intel storage resources. "HPC infrastructure processes large amounts of data very quickly, so it needs fast and reliable storage," says Kiss. GITDA relies on HPE Foundation Care Service for 24x7 support for its HPE servers and storage platforms, helping keep them running and giving GITDA one place to call if there is a problem.

IT manages GITDA's high-performance server infrastructure using the HPE Insight Cluster Management Utility and the HPE Integrated Light-Out Management Engine. "HPE has provided us with a flexible and scalable HPC infrastructure that accommodates the different needs of our users," says Kiss.

Case study
GITDA

Industry
Research

Customer at a glance

Hardware
- HPE ProLiant SL250 Gen8 Servers
- HPE 3PAR StoreServ 7200
- HPE Apollo 8000 System

Software
- HPE Insight Cluster Management Utility
- HPE Integrated Lights Out (iLO)

Services
- HPE Foundation Care

"Our ProLiant SL250 Gen8 Server cluster provides us with a cost-effective solution that delivers high performance, while our Apollo 8000 solution provides us with a system that is smaller in capacity but offers lower operational costs and enables energy-efficient operations."

GITDA has been able to streamline operations and management for its HPC infrastructure, and in the five months that both HPC solutions have been operating GITDA has experienced zero downtime while operating its HPC infrastructure around-the-clock.

"HPE completed both installations on budget, and we now have a reliable, flexible, and scalable HPC infrastructure in place," Kiss states. "We're using GPU-accelerated program code for modeling programs that can be easily accelerated and visualized, and performing the other HPC calculations on our Apollo supercomputers."

GITDA now has over 200 projects running concurrently on its HPC infrastructure, and is evaluating upgrades to improve processing power and capacity. "We are planning to upgrade our HPC infrastructure as demand continues to increase," states Kiss. "For example, we might upgrade our Intel Phi processors to Intel Knights Landing parallel X86 processor to further improve processing speeds. We could also expand our Apollo deployment—which only requires two racks, one of which is used for cooling—to increase the density and energy efficiency of our HPC infrastructure."

He concludes, "It's crucial to offer cutting-edge HPC technology in a country as small as Hungary because it encourages advanced research. This HPC solution from HPE is providing a huge opportunity for our scientists to leverage world-class, advanced infrastructure to achieve .

Learn more at
hpe.com/info/apollo

Sign up for updates

Hewlett Packard Enterprise

intel

a00001461ENW, February 2017