Powering new research, recruiting top faculty

The George Washington University implements a solution that meets researchers' demands for high-performance computing to pursue new, cutting-edge projects

"The Colonial One cluster, based on Dell HPC technology, will tremendously boost the amount of computational power we can dedicate to our projects."

Andrei Alexandru, Assistant Professor, Physics, The George Washington University

Customer profile

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Company	The George Washington University
Industry	Higher Education
Country	United States
Employees	5,000
Website	www.gwu.edu

Challenge

The George Washington University (GW) required new technology to meet demands for high-performance computing.

Solution

The university partnered with Dell™ and Intel to design and deploy a centralized high-performance computing cluster.

Benefits

- High-performance computing solution powers research and delivers results faster
- Researchers can better pursue innovative projects in science and medicine
- GW attracts and retains top scholars
- Centralized cluster saves resources and facilitates interdisciplinary collaboration across the university
- University has a scalable, flexible solution built for the future

Solutions featured

- High-Performance Computing

Located in the heart of the nation's capital, with additional programs in Virginia, The George Washington University (GW) was created by an Act of Congress in 1821. Today, GW is the largest institution of higher education in the District of Columbia. The university offers comprehensive programs of undergraduate and graduate liberal arts study, as well as degree programs in medicine, public health, law, engineering, education, business and international affairs.

> "With the Colonial One solution, I'm hoping that getting results about genetic data is going to take me one or two days instead of a month."
>
> *Diana Lipscomb, Chair, Department of Biological Sciences, The George Washington University*

Each year, GW enrolls a diverse population of undergraduate, graduate and professional students from all 50 states, the District of Columbia, and more than 130 countries. Because GW focuses heavily on innovative scientific research, it is committed to implementing new technologies to support that work. "Our goal is to be the preeminent urban research university," says Sean Connolly, the director of information technology at GW's Columbian College for Arts and Sciences. The university has invested heavily in core networking and data center infrastructure, but historically, those investments were not fully leveraged to support and advance the university's research efforts. To ensure excellence in research, the university sought new technology to meet researchers' demands for high-performance computing solutions. For example, the school's physics department is currently engaged in the study of lattice quantum chromodynamics, a subset of nuclear physics that explores the properties of nuclear particles. Andrei Alexandru, assistant professor in the school's physics department, says, "We put grids on supercomputers, and the grids are the equivalent of solving maybe hundreds of millions of equations with hundreds of millions of unknowns. We require very large computing facilities to be able to carry out this research."

The university's new Computational Biology Institute also needed high-performance computing technology. "In computational biology, there is this insatiable need for computing power," says Keith Crandall, the director of the institute. "In my research, we have a project that is trying to combine all the branches of the tree of life into one gigantic tree. We cannot do this research without high-performance computing." This need for more computing power spanned the entire university. Brian Ensor, the assistant vice president of planning and strategic initiatives at GW, says, "Across the university, from the Columbian College of Arts and Sciences to the School of Medicine, there is a similar need for high-performance computing, data storage, and network capacity for moving large volumes of data."

GW sought to become more attractive to prospective researchers

GW also hoped to enhance the school's reputation as a top research institution in order to attract top faculty talent. "In computational biology and bioinformatics, it's particularly difficult

Technology at work

Hardware
Dell™ NSS high-density storage solution
Dell PowerEdge C8000 Series chassis
Dell PowerVault MD3260 storage array
Dell | Terascala HPC storage solution
Dual Intel® Xeon® processors E5-2670

to recruit faculty," Crandall says. "There aren't many people with the combined skills of computer science and biology, and the few that are out there get snapped up by industry. So to be competitive, we need to have a high-performance computing cluster to allow researchers to do their thing." Diana Lipscomb, the university's chair of biological sciences, echoes that sentiment. "Many of our new faculty are expecting to have computer facilities that they're accustomed to from working at big research institutions," she says.

As GW sought to install high-performance computing technology to meet its needs, it was up against a sizable challenge: a lack of space on campus for computing clusters and a potentially large investment in system administration time. "A number of proposals I saw from researchers were for a small cluster here, and another one there, but we didn't know where we could put those clusters," says Connolly. "As an urban campus, we didn't have the space. Moreover, the task of installing, configuring, and managing these individual clusters throughout their lifecycles is not sustainable, or efficient from an IT staff perspective."

To address this issue, GW focused its efforts on designing a centralized solution. "We started thinking about a more centralized approach, which required potential faculty stakeholders to consider using something collectively," Connolly says.

University implements powerful computing solution in two weeks

The university's search for the right technology led it to Dell. "We are a longstanding Dell customer, and Dell was naturally the partner we looked to in starting the conversation," says Connolly. "We bid this out to four different vendors, and in those conversations no one could provide the value that Dell offered from the partnership perspective. They were partnering with us on the project before they even had our business."

GW collaborated with the Dell High-Performance Computing (HPC) team to design and implement Colonial One, a new shared HPC cluster. "We did a lot of data gathering with our faculty and staff and really partnered with Dell to sit down and design the components to meet our needs as far as density, performance and storage capacity," says Warren Santner, the information systems architect at GW's Columbian College of Arts and Sciences.

Located on the GW Virginia Science and Technology Campus, Colonial One features a 96-node cluster, with all nodes powered by dual Intel® Xeon® processors E5-2670. The cluster is tied together with an FDR InfiniBand network interconnect that features 54.5 Gbps of throughput to each port. This connects all the cluster's nodes to the school's two data storage platforms. The solution is based on the Dell™ PowerEdge C8000 Series, a flexible 4U chassis that houses the cluster's compute, GPU and storage nodes. "Within 4U, we can fit eight compute nodes or four of the double-width sleds to use along with Intel Xeon Phi co-processors, or GPUs, if we choose to go that way," says Santner.

The cluster takes advantage of the Dell NSS high-density storage solution, based on the Dell PowerVault MD3260 storage array, which can store 180 terabytes of data in 4U of rack space. The cluster also includes the Dell | Terascala HPC storage solution, a high-performance parallel file system that provides 300 terabytes of usable capacity. Access to the centralized cluster is open to the entire university community.

The solution was implemented in just two weeks. "The Dell team was very hands-on, working side-by-side with us to help us understand all the moving parts and make sure that we had the right configuration," says Santner. "To go from nothing to a full cluster in two weeks was amazing."

"The Dell team was very hands-on, working side-by-side with us to help us understand all the moving parts and make sure that we had the right configuration. To go from nothing to a full cluster in two weeks was amazing."

Warren Santner, Information Systems Architect, The George Washington University

HPC solution powers research and delivers results faster

GW now has a new high-performance computing tool that gives the university's researchers the compute power they need to do their work. "The Colonial One cluster, based on Dell HPC technology, will tremendously boost the amount of computational power we can dedicate to our projects," says Alexandru. "We can now carry out simulations that are more realistic and closer to physics, which is something we couldn't do before because we didn't have the computational power to do it."

As a result, researchers are expecting to be able to achieve results faster than before. Lipscomb says, "With the Colonial One solution, I'm hoping that getting results about genetic data is going to take me one or two days instead of a month." Adds Crandall: "In my research, we can go from a biological sample like a DNA swab and run the DNA through our software on the high-performance computing solution and have a species and strain identified within 24 hours. That compares with an average turnaround time of one week. And with the faster turnaround time, that means a more responsive treatment. Instead of waiting a week for the most effective treatment, a physician can implement a treatment in a day."

Researchers have a better ability to pursue new projects

University researchers can use the solution's computational power to start new research projects that were not previously possible. For instance, lattice quantum chromodynamics researchers can rely even more heavily on high-performance computing. "With the Dell and Intel-based solution, we have more supercomputing resources available, and new vistas and directions can be researched that would not have been explored otherwise," says Alexandru. "Because of Colonial One, we actually are looking at two or three new problems just because now we have the capabilities to explore them."

GW can attract top faculty members

The university can also bolster its recruitment capabilities by offering prospective instructors and researchers the opportunity to use a state-of-the-art computing solution. "Colonial One is definitely an attractive piece for incoming faculty," Alexandru says. "We had a new faculty member joining this year, and part of his negotiation was a request for computational facilities. He was very happy to learn that he will have resources here for his needs."

Crandall concurs. "This solution is a really effective way to help recruit faculty, because they are looking for these kinds of resources. With the Computational Biology Institute being new to GW, getting exceptional people in the door is going to make the institute. To do that, we need high-performance computing resources, which we now have."

Centralized cluster saves resources and facilitates cross-university collaboration

The centralized high-performance computing cluster will be used to foster new collaboration throughout different departments at GW. "Working with Dell and Intel, we've been able to design a cluster that has a heterogeneous group of components to meet the needs of biologists, physicists, chemists and engineers," says Santner. This shared computing resource will help these different groups save money on research by pooling university resources into a common infrastructure instead of having smaller independent clusters. Ensor adds, "Having this resource helps with the economics of being able to share funding across multiple research efforts."

This collective approach also supports the university's commitment to collaboration. "We have a number of domains coming together for a common need, and that interdisciplinary nature elevates our ability to showcase the great work being done across the university," says Ensor.

University has a scalable, flexible solution built for the future

GW now has a computing solution that has the scalability and flexibility the school needs as it grows. "Our culture at GW has always been to be able to scale what we build," says Connolly. "The modular approach of the Dell C8000 architecture ensures that we can have varying node types, and we can easily scale the architecture. It's built to expand, and we will be able to sustain doubling its size." The Dell technology's storage capacity also enables scalability. Santner says, "With the density in the Dell MD3260 chassis, we can fit 60 drives with 180 terabytes of capacity in a 4U footprint. That will allow us to increase the storage capacity and grow this cluster as the demands on it increase, while maintaining a dense footprint in the data center."

And as Colonial One expands, GW will have a better ability to partner with other institutions. "This isn't just about high-performance computing," Connolly remarks. "We also need to be involved in archiving data and making it available to the general public. As a university initiative, this is helping to define how we fit into the research community in the greater Washington, D.C., area. We can partner with other institutions and make sure that the resources we build at GW are used in a greater context. Overall, this solution fits into a much larger picture for us."

For more information go to www.DellHPCSolutions.com

View all Dell case studies at dell.com/casestudies