



- BACKUP/RECOVERY/ARCHIVING
- CLUSTERING
- CONSOLIDATION
- DATABASE
- MANAGEMENT/UTILITIES
- MESSAGING
- VIRTUALIZATION



CUSTOMER PROFILE

The Flow Physics and Computational Engineering (FPCE) group at Stanford University creates new theories, models, and computational tools for engineering and flow control that advance scientific understanding, leading to better-designed, stronger, lighter, and more efficient products. www.stanford.edu/group/fpc

CHALLENGE

Build a highly available, standards-based high-performance computing (HPC) cluster that supports researchers conducting groundbreaking work in mathematical simulation for engineering design and flow control analysis

SOLUTION

The Dell™ Advanced System Group designed and installed a 212-node, 1,696-processor cluster based on Dell PowerEdge™ servers, knitted together with Cisco Infiniband, Dell PowerConnect™ switches and running the Clustercorp Rocks operating system

BENEFITS

The Dell HPC cluster increases computing power by up to 1,400 percent, enabling significantly larger test cases with higher resolution results, speeding up research, and giving researchers the power to uncover innovative approaches

Engineering by the Numbers

At 15.8 teraFLOPS, Dell cluster boosts compute power by 14X for engineering simulations

The Flow Physics and Computational Engineering (FPCE) group within the Mechanical Engineering Department at Stanford University is staffed with some of the brightest minds tackling one of engineering's biggest challenges. Their mathematical simulations of physical systems analyze phenomena that are inaccessible using traditional experimentation. Modeling the airflow and resistance within an entire jet engine down to the molecular level is just one example. Other areas of research include predicting the dispersion of airborne agents in an urban area and simulating the movement of nutrients in the water surrounding a coral reef. Results are applied in the fields of aerodynamics as well as cooling and material processing. The work of the lab has led to better-designed, stronger, lighter and more efficient products.

Simulations are among the most demanding computational tasks. At Stanford, models working against large data sets can take weeks to complete. Scientists at FPCE never have enough computing power—more is always better. Thus, acquiring the greatest computing power possible within their budget is key to optimizing the results from the group's work.



“The collaborative relationship with Dell plays an important role in our compute model, helping us deploy systems faster that run first-of-kind simulations.”

— Steve Jones

Manager of High-Performance Computing
Flow Physics and Computational Engineering, Stanford University

Dell brings a track record of successful HPC cluster deployments

The existing 48-node, 96-processor cluster at FPCE delivers sufficient computing power for many tasks, but the staff's needs were growing. Pursuing their latest great ideas required an additional, much more powerful cluster.

The Dell PowerEdge server line is well known to Stanford researchers, according to Steve Jones, manager of high-performance computing for FPCE.

Dell servers are highly rated for their remote management and diagnostic capabilities. In fact, Dell engineers helped to define the Intelligent Platform Management Interface standard that Jones relies on for remote diagnostics and management. With the cluster housed across campus from FPCE offices, solving even the smallest problem can become a lengthy jog, so reliability and remote management played prominent roles in the selection process.

The FPCE group had a number of key requirements for the new cluster, including price-performance, the quality of vendor service, support through a single point of contact, availability of replacement parts for the life of the system, and scalability. Relying on his own experience and that of his colleagues, Jones gave Dell top marks across the board.

Deploying an HPC cluster takes experience and coordination. Jones wanted a vendor that could provide a turnkey installation, which would involve integrating products from multiple vendors. After years of HPC leadership, Dell has close relationships with the other vendors that Jones had in mind, including Cisco, American Power Corporation, and Clustercorp. “The strong relationship between Dell and our key partners made Dell an easy choice for us,” he says.

Dell PowerEdge servers deliver 15.8 teraFLOPS of performance

The Dell Advanced Systems Group, which works closely with leading customer and partners on innovative projects such as this, helped Jones design the 212-node, 1,696-processor cluster to meet the group's needs. Compact 1U Dell PowerEdge 1950 servers with two quad-core Intel® Xeon® processors form the basis of the cluster.

Dell PowerEdge servers are a popular choice for HPC clusters. Their price-performance advantage is multiplied in an HPC environment that includes so many servers. Easy to deploy and manage, PowerEdge servers include some of the latest advances in technology while conforming to

industry standards. Several of the world's most powerful clusters are built with PowerEdge servers, including the eighth fastest cluster.¹

High-speed interconnects can make or break cluster performance. For the data communications network, Jones chose 6 Dell PowerConnect 6248 switches, each with 48 ports, for the data transport layer. For message passing, a Cisco 7024 Infiniband switch boosts performance with high-speed nonblocking internode communications. “Every node has a direct connection to the nonblocking switch, which gives us the lowest possible latency,” Jones says. “The compute nodes do not waste cycles waiting for communications, which really unleashes the full processing power of the PowerEdge 1950 servers.”

With so many nodes, an HPC cluster calls for a compact design. Dell PowerConnect 6248 switches have high-performance 10 Gbps interconnects for stacking up to 12 systems. At 48 ports each, a single stack can connect up to 576 servers. Rack density maximizes server connectivity in a 1U form factor. Says Jones, “The PowerConnect switches are very easy to set up and manage.”

Dell Enterprise Deployment Services helps researchers get to work in 11 days

The FPCE group asked Dell to take responsibility for the deployment. In that role, the Dell Enterprise Deployment team coordinated the efforts of all the vendors who took part in the deployment, making sure that every aspect was completed correctly. By not having to divert Stanford staff to work on the deployment, the scientists could continue to focus on their research.

The deployment process began with installation of power and cooling systems. At about five days into that process, the Dell Enterprise Deployment team began to install the 212 servers and 6 switches. “Seeing all that activity happening in parallel was very exciting for our team,” says Jones.

While the Dell team built out the cluster, Jones configured the master node and several compute nodes using the Rocks operating system from Clustercorp—a complete Linux® distribution developed specifically for cluster computing, based on the Rocks Clusters distribution from UCSD. Once all the systems were completely installed, Dell and the Stanford IT team finished deploying the software and then configured and tested the cluster.

Stanford researchers were especially pleased with how quickly they could begin using the new cluster. “We literally began submitting jobs to the cluster on the 11th day, working in full production mode. That's amazing because we have seen other schools take months to deploy an



HPC cluster,” says Jones. “This HPC deployment, managed by Dell, was the fastest I’ve ever seen. This was a large cluster, and having it flawlessly run production scientific code operating with unprecedented fidelity in such a short time is a huge success for our team.”

As Jones reflects back, he notes, “We really had no issues either during or after the deployment. And now, with remote monitoring, the Dell cluster is so efficient to manage that it takes me almost no time. In fact, most of my administrative work is related to application support. As a result of having a smooth running Dell cluster, I can put more effort into designing the next cluster—of course, with the help of the Dell Advanced Systems Group.”

14X improvement in performance helps improve research

Initial tests of the new cluster show it delivers 15.8 teraFLOPS performance, which is 14 times faster than the previous cluster’s 1.1 teraFLOPS. Performance at this level opens the gates to new research. The new cluster returns results in near real time, so that researchers can tweak their calculations more easily and experiment with novel approaches. Researchers get more done, and projects stay on schedule.

Dell helps Stanford create HPC blueprint

The new Dell cluster is a great success. Deployment was completed faster than expected, with no unanticipated issues. The cluster’s high performance is just what researchers were looking for. And the project came in within budget.

HOW IT WORKS

HARDWARE:

- Dell™ PowerEdge™ 1950 servers with quad-core Intel® Xeon® processors
- Dell PowerConnect™ 6248 switches
- Cisco 7024 Infiniband Switch

SOFTWARE:

- Clustercorp Rocks

SERVICES:

- Dell Deployment Services

HPC clusters are in the planning stage at other Stanford departments because they give researchers the power needed to push at the boundaries of knowledge. In fact, the FPCE project was such a success it has become a blueprint for other future HPC clusters. “The collaborative relationship with Dell plays an important role in our compute model, helping us deploy systems faster that run first-of-kind simulations,” says Jones. “We are hoping to see some great new science and groundbreaking work accomplished with Dell HPC clusters at Stanford.”



¹ June 2007 listing, www.top500.org.

September 2007

Printed in the U.S.A.

Dell cannot be responsible for errors in typography, photography, or customer/partner claims made herein. Dell, the Dell logo, PowerConnect, and PowerEdge are trademarks of Dell Inc. Intel, the Intel logo, and Intel Xeon are registered trademarks of the Intel Corporation. Linux is a registered trademark of Linus Torvalds. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others. © 2007 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the written permission of Dell is strictly forbidden.

