

Stanford University Dramatically Expands Compute Capacity



High Performance Computing Center pioneers high-density cluster design that speeds deployments and lowers costs of operation

EXECUTIVE SUMMARY

STANFORD UNIVERSITY HIGH PERFORMANCE COMPUTING CENTER

- Industry: Research, Higher Education
- Location: Palo Alto, California

BUSINESS CHALLENGE

- Provide enough compute power for the in-progress and up-coming scientific research projects.
- Maintain Stanford's position of industry leader, which is required to attract the best talent and funding.
- Work within the physical limitations surrounding space and power on campus.

NETWORK SOLUTION

- Cisco InfiniBand server switches, providing high-speed interconnects within a 212-node cluster of Dell Power Edge processors and Panasas storage solutions.
- Automated software installation and configuration using Clustercorp Rocks™+MOAB

BUSINESS RESULTS

- Stanford achieves high ranking within worldwide HPC community
- "Green" solution sets new standards for energy efficiencies
- Repeatable design delivers long-term cost savings spanning future cluster deployments

Business Challenge

Over the last three years, the High Performance Computing (HPC) Center at Stanford University has grown to include 16 clusters for courses and more than 200 researchers within the School of Engineering. The HPC Center recently completed projects driven by research funded by the National Nuclear Security Administration's Advanced Simulation and Computing (ASC) Program, and Stanford's recently awarded grant from the ASC Predictive Science Academic Alliance Program (PSAAP) is continuing to increase the demand for larger-scale scientific simulations and visualizations.

In anticipation of the PSAAP requirements, the HPC Center began work to significantly increase compute capacity. Steve Jones, HPC manager for Stanford's flow physics and computational engineering group and founder of the HPC Center, said, "We wanted to rapidly deploy new clusters for PSAAP-related research, but we had several limitations that had to be considered. We have always relied on our partners and we have had great success with Cisco in the past, so we turned to them again to help us come up with a scalable, replicable solution."



The biggest challenge was space since Stanford University's campus is at 100% capacity. The new cluster would have to pack compute, interconnects, storage, cooling, and electrical into a small, reclaimed space. Cisco was asked to propose a server interconnect solution that would allow them to maximize the number of connections in each switch, as well as deliver best-in-class performance and scalability. Automation was also paramount – the small staff of the HPC Center required a solution that was simple and almost maintenance free, and the aggressive schedule could not be met without innovative tools for configuring and deploying the cluster.

Network Solution


PSAAP research involves collaborative teams of scientists and their efforts to predict off-design and transient conditions for space travel and hypersonic flight. Simulations make it possible to analyze failure modes before the availability of physical prototypes. The HPC Center wanted to build a new cluster that could handle the much larger compute loads and resolutions required for this analysis. Previous clusters, spanning 48 to 96 processors with 1 gigabyte (GB) of memory each, could only drive visualizations with coarse resolutions. The scale of many simulations required that Stanford's researchers get precious compute time at one of the NNSA Lawrence Livermore, Sandia, or Los Alamos National Laboratories.

The HPC Center aimed to build a new cluster that could give scientists much higher resolution results, and to tackle simulations an order of magnitude larger than previous cluster capacities. "A collaboration between Cisco and Dell gave us an easily deployable solution that could scale to thousands of processors," said Jones. "We chose InfiniBand to optimize performance throughout the cluster. The high-speed interconnects on the Cisco SFS 7000 Series Switches create a full non-blocking fabric that can run jobs with minimal latency. The ability to connect multiple Cisco switches together gives us a great deal of configuration flexibility, which in turn lets us respond very quickly to compute requests."

The Cisco SFS 7012P InfiniBand Server Switches pack 144 10-gigabit per second (Gbps) ports in a 7-rack-unit (7-RU) chassis for a relatively small footprint and extremely efficient power and cooling. Stanford also takes advantage of the 288-port solution, which offers similar efficiencies. The Cisco solution is a perfect complement to the Dell Power Edge compute nodes, which also deliver industry-leading density in terms of computer power per square foot. An innovative "hot aisle" design uses cooled water and variable-speed fans to neutralize temperature increases and allow the compute nodes and InfiniBand switches to completely fill each rack (i.e., 42U of compute nodes in 42U of rack space) and make the most of the precious space and power.

"This Dell and Cisco cluster represents a very cutting-edge and innovative design," said Jones. "Visitors are amazed at the dense rack packing, and the room is cool and quiet – nothing like the wind tunnels of some HPC sites. Our partners have helped us achieve breakthroughs while keeping within the physical limitations of our facilities."





Rapid 11-Day Deployment

The small staff at the HPC Center required a cluster solution that could minimize deployment time as well as simplify ongoing management for the new compute resource. Therefore Cisco and Dell – together with Panasas, the storage solution partner – helped the Center integrate the tools that would yield a high degree of automation. Clustercorp Rocks™+Linux and the MOAB Cluster Suite simplify the software installation across the Dell processors and Cisco InfiniBand switches, and also provide the best workload and resource management utilities availability in the HPC industry today. The total deployment time for the first 212-node cluster was only 11 days, which included electrical and plumbing installations being done in parallel with the set up of the racks and installation of compute, interconnect, and storage components.

“This Dell and Cisco cluster represents a very cutting-edge and innovative design. Our partners have helped us achieve breakthroughs while keeping within the physical limitations of our facilities.”

– Steve Jones, HPC Manager, Flow Physics and Computational Engineering, Stanford University

Business Results

The new cluster design has changed the way that the HPC Center – and other HPC organizations – view cluster installation. The “cookie-cutter” approach has already enabled a second equally impressive cluster deployment, and Stanford will soon go live with a third less than a year after the first. Each of the clusters can be scaled easily by adding in nodes. The Cisco SFS 7000 Series InfiniBand Switches create a fabric that can grow to include thousands of processors, with a single job able to utilize part or all of the processors in the cluster.

With each advancement in cluster computing, the HPC Center helps Stanford foster an environment that attracts the best scientists and instructors. The other milestones achieved with the modular cluster design include:

- **Compute capacity:** The ability to drive higher fidelity or finer-resolution visualizations and massive simulations, which ultimately lead to faster discoveries and insights
- **Density:** The compact processors and interconnects allow far more compute power in a smaller space, which contributes to power and real estate efficiencies
- **Manageability:** Rocks+MOAB simplifies the deployment and administration functions that span processors, interconnects, and storage.

PRODUCT LIST

HPC Server Networking


- Cisco SFX 7012P InfiniBand Server Switches

“Many other groups are looking at our blueprint,” said Jones. “The modularity, automation, and inherent efficiencies of the Cisco and Dell solutions mean that we can rapidly respond to research needs. The new HPC model could be retained even

if we outsource, but today we have gained an in-house resource with a 2-14x improvement in processing power. We can get more done independently, and also be better prepared when we do get access to NNSA clusters.”

For More Information

To find out more about the Cisco InfiniBand server switches, go to: <http://www.clustercorp.com/>.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0805R)