

CUSTOMER SUCCESS STORY

CISCO HPC CASE STUDY: NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS (NCSA) BUILDS OIL AND GAS SEISMIC EXPLORATION CLUSTER

The demand for dependable compute power for scientific research continues to grow as ever more complex problems are tackled and high-performance compute resources become more and more critical as tools for success. Together, Cisco® and the NCSA have built a reliable, high-performance cluster for scientific exploration, which is based on the industry-standard building blocks of the Cisco InfiniBand solution, Intel EM64T, and Linux.

Figure 1 shows how each compute server has one InfiniBand connection for computing and one Ethernet connection for management.

Figure 1. Compute Server Connections




Figure 2 shows how the overall cluster is built from multiple racks of servers in the NCSA data center.

Figure 2. Cluster Built from Multiple Server Racks



EXECUTIVE SUMMARY

Customer Name

National Center for Supercomputing Applications (NCSA)

Industry

High-Performance Computing, Education, and Research

Business Challenges:

- Build a flexible and powerful high-performance computing platform that empowers scientists and engineers in many different scientific disciplines and industries.
- Provide an innovative solution based on a industry standard server building block. Open standards to maximize price-performance based on dollar investment and provide compatibility for customer applications.
- Contribute to the future of computing by researching and deploying an innovative system that decreases the cost and/or extends the range of computational science and engineering.
- Use a network design that allows for future performance upgrades with no architecture changes and minimal system disruption.
- Provide rapid deployment and confidence in the system. Continue meeting requirements of customers, who, like the NCSA, are contractually obligated to bring the system up and have it fully operational in a short period of time.

Network Solution

A high-performance InfiniBand server fabric with Cisco SFS 7000 and 7008 InfiniBand server switches, Cisco InfiniBand host channel adapters, and Cisco InfiniBand host drivers and scalable fabric management.

Business Value

- Deployed an innovative high-performance supercomputing platform based on industry standard Dell EM64T servers and ultralow-latency InfiniBand. The resulting system provides higher performance than a similar size proprietary system at the NCSA, but for significantly less cost.
- The solution is completely based on open standards that are compatible with the NCSA's customers' needs, which allows the NCSA to draw from and contribute to the open source community.
- The NCSA was able to meet its contractual obligations. The supercomputer was deployed, debugged, stable, and running customer codes within two weeks of purchase order issuance.
- The solution is both a functional solution to the customers' needs and a showcase for the future of computing. The NCSA pushed the boundaries of what was known about InfiniBand clustering with this system.
- The design allows for an easy upgrade path to increase network performance with no redesign and minimal disruption to the existing supercomputer.
- The system has been running in production at nearly full utilization since its deployment over six months ago, with no unscheduled system downtime.

“NCSA has experienced unprecedented reliability, performance, and support with the Cisco InfiniBand solution,” said NCSA Senior Operations Manager Brian Kucic. “Cisco has proven InfiniBand is ready for research and commercial HPC.”

— Brian Kucic, Senior Operations Manager, National Center for Supercomputing Applications

BUSINESS CHALLENGE

The NCSA at the University of Illinois at Urbana-Champaign has two decades of experience providing high-performance computing resources to scientists, engineers, and the private sector. The NCSA has earned and maintains the reputation of being an innovative center for new technology, pushing the bounds of computing, networking, storage, data mining, and visualization. The NCSA, along with the San Diego Supercomputer Center and Pittsburgh Supercomputing Center, is one of three centers supported by the National Science Foundation with the mission statement of research, discovery, and education. The NCSA's three main principles are:

- Enabling discovery at the leading edge by providing advanced cyberresources
- Empowering all scientists and engineers through cyber environments that allow ready access to these advanced computing resources
- Realizing the future of computing by researching and deploying innovative systems that decrease the cost and/or extend the range of computational science and engineering

Scientists and engineers must have computing systems that are accessible, robust, and easy to use to advance scientific discovery and the state of the art in engineering. Through close collaboration among vendors, NCSA staff, and the research community, the NCSA provides platforms for frontier science and engineering. One of the organization's core missions is to expand the affordability and capabilities of scientific computing and cyberinfrastructure to both research and commercial computing environments.

PROBLEM AT HAND

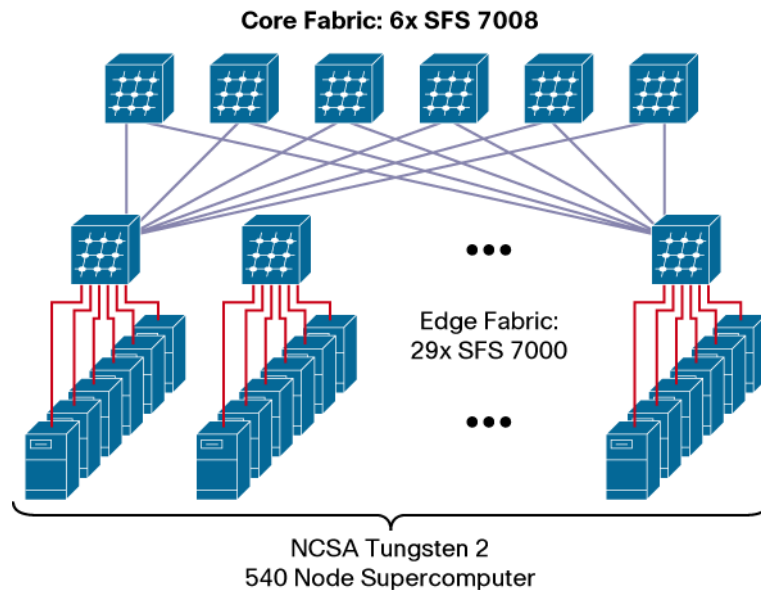
The NCSA needed to increase its compute infrastructure in a short timeframe because it was under contract with a commercial company in the oil and gas space. This commercial company needed to have a powerful platform that would be capable of running demanding parallel seismic exploration codes. The time to create solutions for these applications is very important for this company's core business. However, the NCSA had a fixed budget within which to work, so it needed a powerful machine that was designed with cost-effective and reliable industry-standard parts. The machine also had to be easy to run and manage and be extremely reliable, because the NCSA was not adding extra staff for the new cluster.

The research being done on this cluster is critical to the success of the NCSA's customers' business. The NCSA took this on as a research challenge to find a solution to its problem.

BUSINESS SOLUTION

Weighing its requirements, the NCSA decided on a Linux-based industry standard server cluster, with the latest Intel-based processors available. Based on its budget and the desire to have a very high-performance, reliable interconnect, the NCSA selected Cisco InfiniBand as the server cluster interconnect. The NCSA furthered the advancement of supercomputing by showing the performance that could be derived from clusters made of InfiniBand and Intel EM64T processors. Based on the performance requirements, the cluster size of 540 nodes was selected. (See Figure 3.)

Figure 3. NCSA Server Cluster Model



Two weeks after the purchase was finalized, the Tungsten 2 (T2) supercomputer was born. T2 is a 540-node InfiniBand cluster based on Cisco InfiniBand switching technology and Dell PowerEdge 1850 servers, which are equipped with Intel EM64T 3.6-GHz dual processors. The InfiniBand fabric design is based on a two-tier Clos-style network with edge switches connecting the hosts and core switches comprising the backbone of the fabric. This design minimizes latency in the fabric without compromising any bandwidth. The InfiniBand interconnect linking the nodes can transfer 800 gigabits of data per second (Gbps), with less than a 6-microsecond average delay in the point-to-point transmission of data. This high-speed data transfer enables users to run tightly coupled applications that run at highly optimized performance efficiencies.

Impressively, this cluster was able to be brought up to full service in a very short time, enabling the customer to start research right away on the cluster. Working together, Dell and Cisco were able to deploy and bring up the cluster, debug problems, and hand the cluster over to the customer in about two weeks. It took another week to debug some application interaction issues, but the cluster has been running perfectly ever since. The Cisco InfiniBand fabric has not had any major problems since the cluster was brought up. T2 is still currently utilized at near maximum capacity by corporate researchers in the oil and gas industry.

WHY CISCO?

Cisco provided a well-tested solution that combines state-of-the-art InfiniBand clustering technology, fabric management, server adapters, and upper layer protocols. Cisco also provided extensive design tools, onsite bringup and tuning capabilities, and world-class service and support. Because of the extensive Cisco experience with InfiniBand and high-performance clusters, Cisco was able to help the NCSA bring up T2 and stabilize it quickly, getting the cluster into a production environment in a matter of weeks. Cisco and Dell were able to deliver a very solid platform and, with the NCSA, to bring the cluster up more quickly than any InfiniBand fabric of a similar size has previously been deployed.

To date, the cluster has been running at expected performance levels with no unexpected problems with the Cisco SFS InfiniBand fabric.

NEXT STEPS

The NCSA is very happy with the solution and the performance of the Cisco InfiniBand server fabric of T2. In addition to upgrading the existing T2 InfiniBand fabric, the NCSA is also investigating InfiniBand fabrics on a much larger scale than 540 nodes.

FOR MORE INFORMATION

To find out more about Cisco SFS solutions, visit <http://www.cisco.com/en/US/products/ps6418/index.html>.

To find out more about the NCSA, visit <http://www.ncsa.uiuc.edu>.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath 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. (0502R) C36-377865-00 11/06

