

Accelerating Physics Research with High-Performance Global Networks

High-speed routers and switches help Caltech and global research partners handle data from huge physics experiments at the frontier of science.

EXECUTIVE SUMMARY
<p>CALIFORNIA INSTITUTE OF TECHNOLOGY (CALTECH)</p> <ul style="list-style-type: none"> • Education and research • Pasadena, California • Over 2600 <p>BUSINESS CHALLENGE</p> <ul style="list-style-type: none"> • Prepare for startup of major physics experiments, CMS and ATLAS, at the Large Hadron Collider at CERN (Geneva) in 2007 • Overcome limitations of long-range networks and storage to share massive amounts of data among global research collaborations • Add capacity and features while protecting existing investments in equipment <p>SOLUTION</p> <ul style="list-style-type: none"> • High-speed intelligent switched IP network with 10 Gigabit Ethernet line cards • MPLS, VLANs, and policy-based routing features help segment, prioritize, and direct traffic flows across the global computing grid <p>BUSINESS RESULTS</p> <ul style="list-style-type: none"> • Upgraded to higher density switches without changing switching chassis • Achieved high performance in a production network to share, process, and analyze many terabytes of physics data. • Maintained flexibility to mix and match small- and medium-sized servers while fully utilizing high-speed bandwidth

Business Challenge

In 2007, the world's largest and most powerful particle accelerator is scheduled to start up at the CERN physics center in Geneva, Switzerland. The Large Hadron Collider (LHC) is capable of producing hundreds of millions of collisions per second, reproducing conditions in the universe just a few million-millionths of a second after the Big Bang. The LHC's enormous concentration of energy and unprecedented particle-collision rate will enable researchers around the world to study fundamental particle interactions that may lead to discoveries such as the origin of particle mass and the nature of the mysterious dark matter in the universe.

The California Institute of Technology (Caltech) and its research partners have developed the Compact Muon Solenoid (CMS), a 12,000-ton ensemble of state-of-the-art particle detectors housed in the most ambitious superconducting magnet ever built. The CMS magnet will record the energy, charge, and tracks of particles emerging from proton-proton collisions. The CMS experiment, one of the two

largest of five LHC-related experiments, is a collaboration of more than 2,000 physicists and engineers from 182 institutions in 38 countries.

The key to discoveries is the ability of scientists to repeatedly access terabyte-scale and larger data samples from the LHC and seek the rare signals of new physics in the avalanche of already-understood particle interactions.

Caltech physics professor Harvey Newman is the board chair for the U.S. CMS Collaboration and also heads Caltech's UltraLight project. The UltraLight project has developed a global research grid based on a Cisco® high-performance network, utilizing parts of the national facility provided by Cisco on the National LambdaRail backbone. The network interlinks the CMS institutions and individual research teams to allow shared access to multiple terabytes of data at unprecedented speeds.

According to Newman, “Over the last few years, our team and our collaborators have addressed, and largely overcome, many of the challenges of disk- to-disk transfer, including TCP protocol, disk buffering, network interface, and Linux kernel. Solving the challenge of disk-to-disk transfer, which is now only constrained by the read-write speed of disks, is crucial for distributing as much data as possible, as quickly as possible, across the globally distributed grid system.”

Network Solution

The existing network designed by the Ultralight project team is based on Cisco Catalyst® 6500 Series Switches and Cisco 7600 Series Routers. In 2005 this Caltech-led team won the Supercomputing Bandwidth Challenge for the third straight year, transferring physics data at a rate of over 150 Gbps over the Cisco network, equivalent to downloading more than 130 DVD movies in one minute. The team also developed a “hyper-challenge” for Supercomputing 2006 in cooperation with Cisco, once again using fully populated Catalyst 6509-E switches to support storage-to-storage transfers at full speed to and from large sets of disks located at CERN, in Brazil, and in South Korea.

The network supports a distributed computing environment in which data from the Tier-0 CERN laboratory is sent to Tier-1 national research centers around the world. These centers, which are capable of storing multiple petabytes of data, then distribute data samples to university-based Tier-2 centers where physicists perform most of the required data analysis and simulation.

A Tier-1 data center has 500 to 1000 nodes; each Tier-2 data center has approximately 100-200 nodes. All nodes have four disks, and by 2008 most nodes will have 1-terabyte (1000 GB) disks, or 4 TB of disk space in each node. “We needed to find a way to overcome the limitations of these slow electromechanical devices by increasing switching capacity. This would allow more servers to transfer data over the network at the same time,” says Newman. “We also hoped to add this

“We are on track to make use of our networks in support of our science, limited only by the read and write speeds of disks, not the networks.”

— Harvey Newman, Professor of Physics, Caltech

capacity while preserving our investment in the existing network.”

The Cisco Catalyst 6509-E Switches were equipped with Cisco Catalyst 6704 4-port, 10-Gbps Ethernet line cards, providing a density of 32 10-Gbps ports per chassis. Newman and his team were among the first to install the new Cisco Catalyst 6708 8-port 10-Gbps Ethernet module, doubling the port capacity of the switch. “By increasing the throughput of the network using the higher density switch, we could transfer data from more disks at the same time to compensate for the read-write speed limitations,” says Newman. “The challenge is to refresh 200 terabytes of disk cache without stressing the network. We pushed the 6708 to the limit with all channels loaded, and we were able to achieve multiple ten gigabits per second between the Pasadena campus and a point of presence in Los Angeles. The Cisco equipment performed admirably.”

In addition, about half of the servers in the grid have 10-Gbps Ethernet interfaces and the rest have 1-Gbps Ethernet ports. “Our application does not require fielding very large disk servers,” says Newman. “It is more cost-effective for the global grid to use a mix of small and some medium-sized servers. The 6708 module, along with 6748 48-port 10-Gbps Ethernet modules, also gives us the room and flexibility to support a mix of 1-Gbps and 10-Gbps server interfaces in the same chassis. We could concentrate more of these cost-effective small and medium-sized servers to use our high-speed bandwidth fully.”

Several projects are being integrated through the UltraLight initiative. The team is using multiple intelligent service features on the switches and routers, including VLANs to segment project traffic, IP Multiprotocol Label Switching (MPLS) to facilitate interconnectivity with different locations, and policy-based routing, to build Layer 1, 2, and 3 paths dynamically. “We are enabling a lot of different features and we haven’t experienced any conflicts or performance degradation in the network,” says Newman. “With the Cisco equipment, we did not have to sacrifice functionality to get the performance we needed.”

Business Results

There are six Cisco Catalyst 6509 Switches located at the university-based centers represented by Caltech, the University of Michigan, and the University of Florida, as well as the Stanford Linear Accelerator Center and Fermilab. According to Newman, the switch provided a “natural evolution to higher throughput by allowing us to replace the 4-port 10-Gbps modules with the 8-port 10-Gbps modules. We are able to preserve our investment in the existing equipment, while adding substantial throughput to the network.”

That investment protection also extends to the servers. The Cisco switches provide the flexibility to mix and match interfaces, enabling the researchers in Tier-1 and Tier-2 centers to support and upgrade servers more easily to keep pace with data-intensive computing demands.

“The report card is extremely good. The Cisco switches just operate and we don’t have to worry about them,” says Newman. “We are on track to make use of our networks in support of our science, limited only by the read and write speeds of disks, not the networks.”

PRODUCT LIST

Routing and Switching

- Cisco Catalyst 6500 Series Switches
- Cisco 7600 Routers

NEXT STEPS

“The network will revolutionize data-intensive grid computing and other forms of scientific computing over the next decade,” says Newman. “It is an enabling force in the next round of scientific

discoveries expected when the LHC begins operation.” Newman also expects that many of UltraLight’s developments in the areas of networking, monitoring, management, and collaborative research to be applicable to many fields of data intensive e-science.

FOR MORE INFORMATION

To find out more about Cisco routing solutions, visit: <http://www.cisco.com/go/routing>

To find out more about Cisco switching solutions, visit: <http://www.cisco.com/go/switching>

To find out more about Caltech, visit: <http://www.caltech.edu>

This customer story is based on information provided by Professor Harvey Newman at The California Institute of Technology and describes how his organization benefits from the deployment of Cisco products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.



Americas 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: 408 527-0883

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

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

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.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, 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. (0609R)