DEPARTMENT OF ENERGY'S NETWORK BREAKS MASSIVE DATA TRANSFER BARRIERS TO UNLOCK UNIVERSE MYSTERIES

Summary

Industry: Scientific research

Challenge:

- Overcome inherent difficulties of transmitting high-volume, bulk data transfers
- Offer guaranteed bandwidth service and traffic engineering across multiple domains to support analysis of LHC data flows

Selection Criteria: ESnet needed a 10Gbps ethernet routing solution for its core network with a flexible architecture that supported IPv6, MPLS, and ESnet's reservation system. It needed to have carrier-grade reliability and high-availability

Network Solution:

- Juniper Networks MX960 and MX480 Ethernet Services Routers
- Juniper Networks EX4200 line of ethernet switches

Results:

- Provided reliable, high-speed, connectivity between U.S. universities, labs and international research institutions required to support the inherently collaborative, global nature of modern, large-scale science
- Supported advanced MPLS capabilities and ESnet's OSCARS on-demand bandwidth reservation system
- Upgraded network infrastructure from OC192 to 10 Gbps Ethernet in less than five months

ESnet
Energy Sciences Network

The European Organization for Nuclear Research's (CERN's) Large Hadron Collider (LHC) will allow physicists to ponder the mysteries of the universe and examine matter as it existed just after the Big Bang. But the scientific discoveries from these subatomic smashups are predicated on the ability to easily distribute massive data sets from the LHC to scientists around the world. In the United States, thousands of researchers at national laboratories, universities, and other institutions use Energy Sciences Network (ESnet) to communicate and collaborate on the LHC projects, as well as many other scientific challenges funded by the Department of Energy (DOE). ESnet, which is funded by the DOE Office of Science and managed by Lawrence Berkeley National Laboratory, provides direct high-performance connections to more than 40 major DOE sites as well as fast connections to more than 100 other networks.

Challenges

ESnet needed enough bandwidth to transport multiple 10 Gbps information streams—the equivalent of transmitting 500 hours of digital music per second for each 10 Gbps line. That called for an upgrade from an OC192 network to a 10 Gbps Ethernet infrastructure. By 2010, the LHC is projected to move 100 Gbps of data nonstop and ESnet will be the sole conduit for U.S. researchers accessing that data. "We had to upgrade the network to support the existing gigabytes of traffic, in addition to the data coming from the LHC," said Kevin Oberman, senior network engineer at ESnet.

The need for high performance doesn't stop with the LHC. "Big science has discovered that networks are an effective way to collaborate," said Oberman. "In the past, they've assumed that networks couldn't handle the load and they transferred large data sets on DVDs or USBs. More big science projects will use our network, which will place more demands on ESnet. We've only begun to grow, and we're going to get a lot bigger."

Selection Criteria

ESnet needed a routing solution that would support 10 Gbps Ethernet at the core. High performance was critical, as was a flexible architecture that supported advanced services, including IPv6 and MPLS. The ideal router also had to deliver carrier-grade reliability and high availability.

"We needed a routing platform that fit our budget and would do the job," said Oberman. When Juniper announced the MX960 Ethernet Services Router, Oberman instantly liked what he heard. "The MX960 had all of the capabilities, hardware interfaces, and optics that we needed," Oberman said.

ESnet also needed Ethernet access switches to connect its major hub locations to the core. Based on the robustness and capabilities of the then newly announced Juniper Networks EX Series Ethernet Switches, ESnet selected the new equipment.

Solution

ESnet deployed the Juniper Networks MX Series Ethernet Services Routers in the core 10 Gbps Ethernet network that connects the DOE facilities and collaborators. It also uses EX Series Ethernet switches to provide 1 Gbps connectivity to labs and universities responsible for analyzing the LHC data. As a Tier 1 service provider, ESnet also connects to Internet2 and has multiple international peering points (see figure).

The Juniper Networks MX Series Ethernet Services Router is a seamless, cost-effective way to deploy Ethernet, and is designed to meet the sophisticated requirements of advanced services and applications. The MX Series delivers performance at scale, advanced quality of service (QoS), and service flexibility. Dense Port Concentrator (DPC) cards provide multiple physical interfaces and Packet Forwarding Engines on a single board. The MX Series features superior QoS at the interface level and per VLAN to ensure that services receive the appropriate level of quality regardless of traffic conditions. The MX Series leverages Junos Software, giving the MX Series richness, stability, and service breadth. The MX960 router is one of the industry's largest capacity carrier Ethernet platforms with up to 960 Gbps of switching and routing capacity; the MX480 router provides up to 480 Gbps of capacity.

The Juniper Networks EX4200 line with Virtual Chassis technology provides high-performance, carrier-class ethernet switches that are easy to deploy and manage, allowing for more energy efficient networks that reduce capital and operational expenses. They are also powered by the field-proven Junos Software, which leads to operational efficiency across the switch and router products. All EX4200 line switches include HA features such as redundant, hot-swappable internal power supplies, and field-replaceable, multi-blower fan trays to ensure maximum uptime.

Researchers analyzing the LHC data reserve bandwidth for the data transfers as they need it. Fermilab and Brookhaven National Laboratory are the "Tier 1" centers in the U.S., making the data available to many other universities and labs where researchers will analyze the data sets. Data streams from the LHC are often in excess of 1 Gbps of bandwidth, and transferring such massive amounts of data within set time constraints and across multiple domains can be difficult because of network latency. ESnet is exploring the use of dynamically provisioned QoS paths as the solution. The On-Demand Secure Circuits and Advance Reservation System (OSCARS) facilitates the on-demand provisioning of guaranteed bandwidth secure circuits within ESnet. Using OSCARS, researchers can reserve a specified amount of bandwidth in the 10 Gbps path, and any other traffic is marked as best-effort or scavenger.

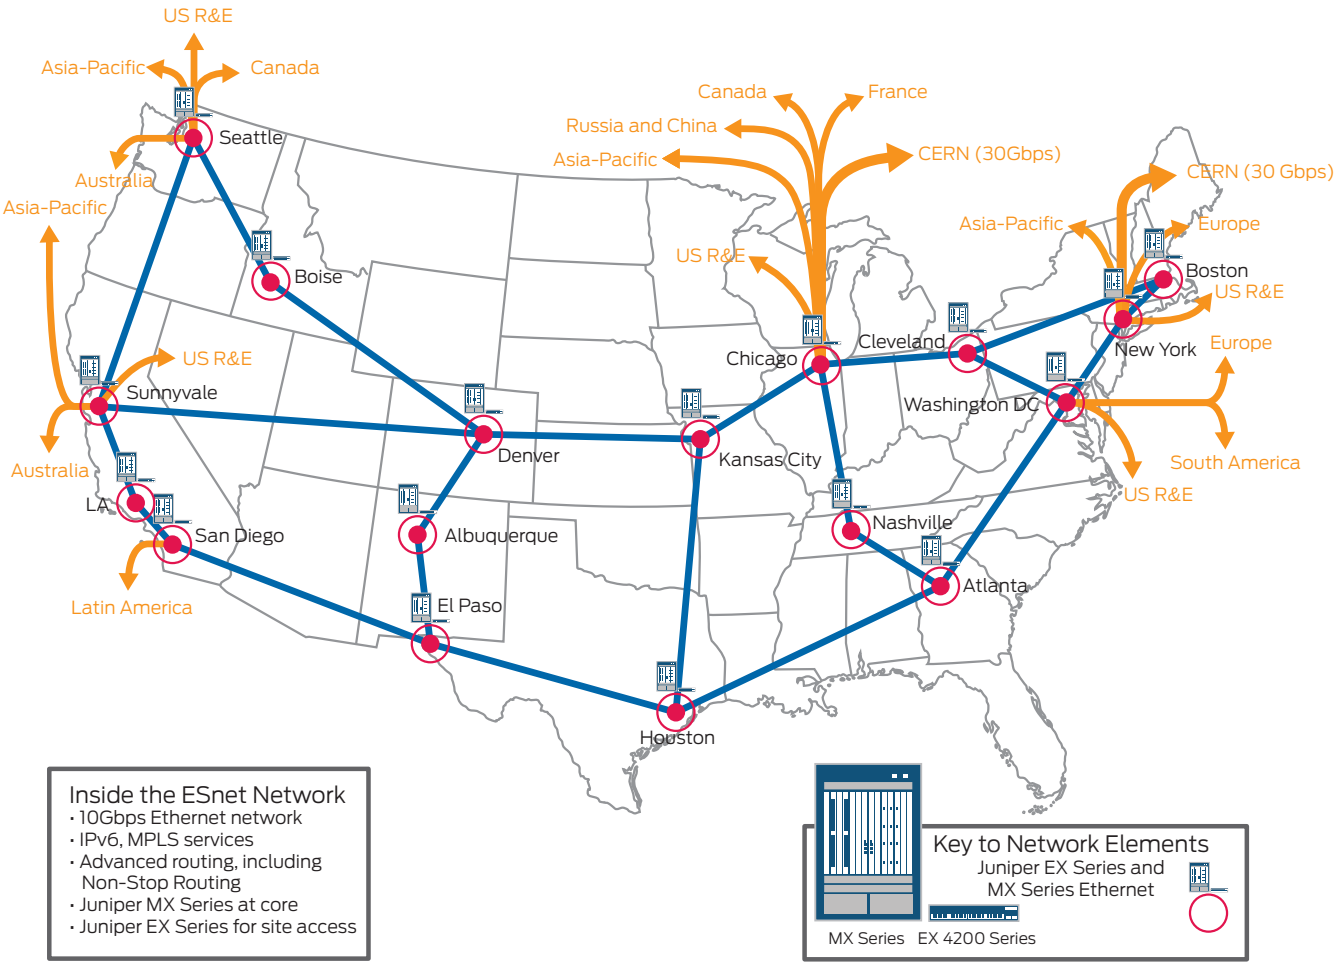

"The way we're doing MPLS with pseudo-wires is a new approach, and most router vendors didn't support it," said Eli Dart, senior network engineer at ESnet. "Juniper has a pseudo-wire capability that met our needs nicely and greatly simplified our software. It saved us significant development time. We're not absolutely convinced that we could have implemented this capability at all if we didn't have exactly what we wanted."

The MX Series routers enabled ESnet to deliver high levels of service availability. "Our goal is to achieve five 9s of availability, and we're not happy if we don't comfortably exceed four 9s," said Oberman.

The MX Series routers deliver carrier-grade reliability, including HA features that ensure that the network is always running, including Non-Stop Routing (NSR), MPLS fast reroute, and unified in-service software upgrade (ISSU). The MX Series routers have common hardware redundancy including the Switch Control Board, Routing Engines, fan trays, and multiple power supplies. Additional MX Series routers are also standing by at hub sites. "The high availability and recovery capabilities that Junos provides are really important," Oberman said.

Many of the ESnet locations use the NSR feature, which allows a router with redundant Routing Engines to switch over from a primary Routing Engine to a backup Routing Engine without alerting peer nodes that a change has occurred. With ISSU, ESnet can upgrade between different versions of Junos Software without incurring network downtime, which ultimately delivers a higher service level at a lower cost.

ESnet's expectations and requirements have been met by the EX4200 line of switches, which provide 1 Gbps access to labs and universities responsible for analyzing the LHC data "The EX Series Ethernet Switch gives us the ability to plug in optics or copper, so we can run at 100 Mbps or 10 Gbps speeds as necessary," said Oberman. Most locations have one EX Series Ethernet Switch, but as performance needs grow, ESnet can easily grow incrementally and also take advantage of Virtual Chassis technology to interconnect up to 10 EX4200 switches, and manage them as a single, logical device—delivering the reliability, scalability, and manageability of traditional chassis-based systems.

"We have hubs scattered all over the country," said Oberman. "If we need to replace a switch, we can't just run downstairs to the telco room. We like that the switches run Junos Software, and the cost/performance of the switches is very good."

"We have a unique customer set that requires the ability to handle large data flows. We've engineered our network infrastructure to meet the requirements of large science, and the Juniper MX Series has allowed us to create a highly reliable IP infrastructure that is cost-effective, flexible and supports the software we've developed."

Steve Cotter,
Head of ESnet

Operational efficiency is equally important to ESnet. With Junos running across its core router and switch platforms, ESnet's network engineering staff can quickly and efficiently keep up with changing business demands. ESnet is a long time user of Juniper routers, which made for a swift rollout of the 10 Gbps Ethernet network.

"The rollout was amazingly smooth, with no downtime seen in the backbone. We were very pleased, and having Junos across all platforms certainly saved us a lot of time," said Oberman. "The time it takes to translate configurations is not trivial, and we have limited engineering manpower. Because of the design of Junos, it is much faster to write a configuration in Junos than any other network operating system".

Next Steps and Lessons Learned

Supporting the massive information from the LHC is just one large-scale science project that is currently underway. ESnet's network engineering team is preparing to support several large telescope projects that will bring in huge amounts of astrophysics data in the coming years. "We've only begun to grow, and we're only going to get bigger," said Oberman.

For More Information

To find out more about Juniper Networks products and solutions, visit www.juniper.net.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net.**

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

3520290-002-EN Oct 2009

Printed on recycled paper