Customer PROFILE

Smashing Atoms: CERN turns to Force10 To Support Next-generation Physics Research

Customer
CERN

Industry
Scientific Research

Applications
Grid Computing

Highlights
The Force10 TeraScale E-Series switch/routers meet CERN's demanding requirements for their global network supporting the world's largest scientific instrument. The E-Series provides 10 Gigabit Ethernet (GbE) LAN and WAN interfaces, non-blocking, high density equipment, fault tolerance, adherence to standards, and a low cost of ownership.

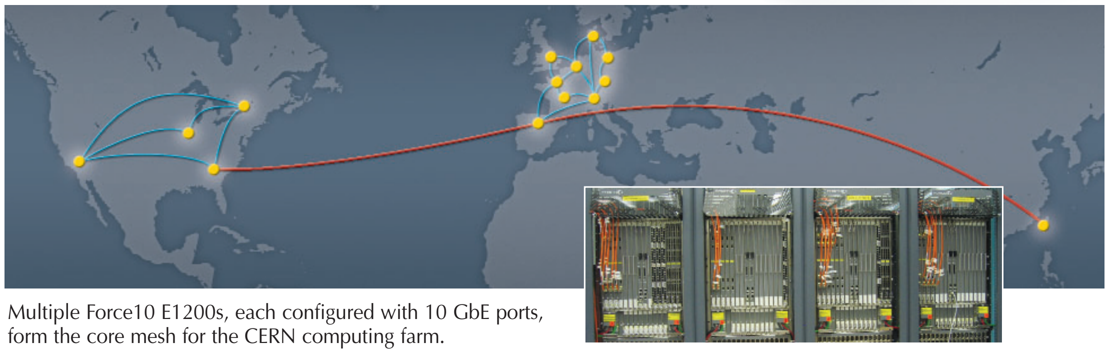

Multiple Force10 E1200s, each configured with 10 GbE ports, form the core mesh for the CERN computing farm.

How do you take data from the world's largest proton collider and distribute it to scientists around the globe? That's the challenge the technology group at Geneva-based CERN faced in designing the network to support the Large Hadron Collider (LHC), which is slated to go on-line in 2008.

The LHC, which is being built in the same tunnels as CERN's 27 kilometer long Large Electron Positron collider, will accelerate proton beams, producing an estimated 15 petabytes (15 million Gigabytes) of data annually. Huge detectors will capture what happens when the beams collide and feed the data into a computing grid for initial processing. From CERN the data will be distributed to a series of Tier-1 computing centers for subsequent analysis and distribution, providing almost 7000 scientists in around 500 institutes and universities with data from the LHC experiments.

"Networking is a highly important part of the overall infrastructure," says David Foster, communications, systems and networking group leader at CERN. "It touches the experiments themselves, the computing center at CERN, and it supports distributing the data out to the collaborating institutes."

The members of the LHC computing team are building the network from the ground up to interconnect the LHC computer center's farm of 6,000 machines and to support 100 Gigabits per second (Gbps) of aggregate bandwidth out of the laboratory to Tier-1 centers in Europe, North America, and Asia.

In addition, the team needed network infrastructure to support the various LHC experiments; one experiment alone requires 1,000 one-Gbps ports. Given the scale of the new infrastructure it's installing, the team also decided to upgrade the lab's campus backbone to a 600 Gbps core.

Stiff Requirements

Designing a network to support the world's largest scientific instrument translates to a demanding list of requirements that networking gear must meet. For starters, the LHC team wanted to build a non-blocking router core capable of aggregate bandwidth of about 2.4 terabits per second, according to Foster. In addition to high speed, non-blocking devices with high port density, the LHC computing team also required fault tolerant devices that adhere to standards.

"Fault tolerance is a major requirement for us," he says. "We need that in two dimensions — in the product itself and in the architecture we put in place." Along with detailed technical specifications, the CERN networking team factored into its requirements the total cost of ownership over the life cycle of the networking equipment. Then they began soliciting bids from a variety of vendors, including Force10 Networks.

"We are always looking for state of the art products because our requirements are on the limits of what is commercially available," Foster says. "So we're aware of the vendors that are pushing the state of the art, and Force10 is well known for being a leading light in the 10 Gig Ethernet space."

Smashing Atoms: CERN turns to Force10 To Support Next-generation Physics Research

Customer PROFILE

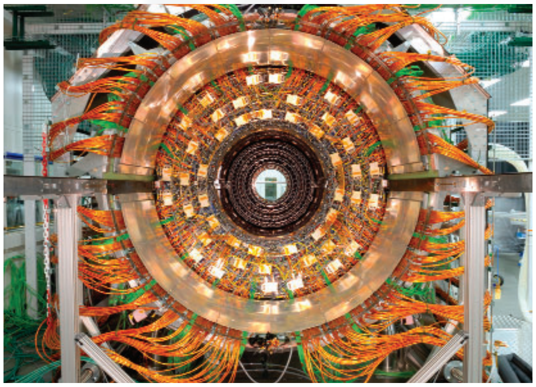

"Force10 is well known for being a leading light in the 10 Gig Ethernet space."

David Foster
Communications, Systems and Networking Group Leader, CERN

Force10's TeraScale E-Series switch/routers met all of the networking team's requirements, Foster says, providing 10 Gigabit Ethernet (GbE) LAN and WAN interfaces, non-blocking, high density equipment, fault tolerance, adherence to standards, and a low cost of ownership. "Force10 delivers carrier-class products at enterprise prices," he notes.

Meeting the Needs

The networking team is in the process of deploying a combination of Force10 E1200s and E600s. For example, both classes of switch/router will be used to support the experiments themselves. In addition, eight E1200s, each configured with 40 10 GbE ports, will form the core mesh for the computing farm and connect to an aggregation layer consisting of two more E1200s. To support connectivity to collaborating Tier-1 research centers, the team has configured two E600s with 10 GbE WAN interfaces, which provide transatlantic connections to sites in the United States. These E600s are linked, in turn, to two E1200s which tie in the other Tier 1 sites.

Each E1200 has a 1.68 Tbps non-blocking switch fabric, a forwarding capacity of 1 billion packets per second, and 14 linecard slots which can be configured with a mix of up to 672 GbE or 56 10 GbE line-rate, non-blocking ports. Each E600 has a 900-Gbps non-blocking switch fabric, a forwarding capacity of 500 million packets per second, and seven linecard slots which can support a mix of up to 336 GbE or 28 10 GbE line-rate ports.

Designed for carrier-class switching and routing, the E-Series has superior performance, redundancy and availability features. For example, the E-Series linecards, switch fabric, backplane, central processor and operating system have all been optimized to process terabits of traffic at line rate in a reliable, predictable fashion.

To maximize up-time, the E-Series has a patented three CPU architecture, with switching, routing and management functions running on three distinct processors. This fully distributed, multiprocessor design protects each control plane process, allowing a fault in one control plane to be contained while protecting other parts of the system. In addition, all E-Series devices feature fully redundant components. In the event of a failure, secondary components continue to process traffic with no packet loss, ensuring hitless failover.

The CERN networking team is taking advantage of all of the E-Series' redundancy features, including dual supervisors and switching fabric, Foster says. "Force10 was also competitive because it has a very good software implementation that adheres to the standards we were looking for."

Force10 has robust implementations of BGP, IS-IS, OSPF and RIP routing protocols as well as Rapid Spanning Tree for optimal switching. By supporting active redundant links, Force10 provides immediate, zero-packet loss failover if one link fails. In addition to VRRP for router redundancy, Force10 supports graceful restart of BGP and OSPF and the VLAN redundancy protocol.

FORCE10

Smashing Atoms: CERN turns to Force10 To Support Next-generation Physics Research

Customer PROFILE

Beyond Dollars and Cents

CERN anticipates both short-term and long-term benefits from its partnership with Force10. "We're looking at a 20x performance gain over our current infrastructure," says Foster. He also expects cost savings from the new network.

CERN's ability to use Force10's 10 GbE WAN PHY interfaces over long-haul links has enabled the lab to move away from SONET for its point-to-point transatlantic links, "so we're enjoying some savings there," Foster says. In addition, using Force10 gear throughout the LHC computing build-out will simplify ongoing operations and support.

"We're able to use the Force10 equipment in all the aspects of the network infrastructure, from the campus networking to the farm to the experiments to the wide area networking," says Foster. "That level of homogeneity should reduce costs in terms of complexity."

The scalability of the Force10 E-Series will also allow for non-disruptive growth. "The scalability is very good and that's a very interesting feature for us. We have a farm of 6,000 machines, but if we need to add 1,000 machines, we can scale up the network without having to re-engineer everything," Foster says.

One benefit that's hard to put a price tag on is the technical support Force10 is providing. "As with any deployment of this type there are always, if not surprises, things to be resolved. We're pushing the boxes to the limit," Foster notes. "Since the very beginning we've had extremely good technical contact with Force10. This is very, very important to us. We're very fortunate to have world class engineers here at CERN and so consequently our conversations with vendors, such as Force10, are at an extremely high level."

He adds: "The excellent technical relationship we have with Force10 is very important in that we can address problems quickly and deliver networking solutions to our researchers with confidence and on time."

"We have a farm of 6,000 machines, but if we need to add 1,000 machines, we can scale up the network without having to re-engineer everything."

David Foster
Communications, Systems and Networking Group Leader, CERN

Force10 Networks, Inc.
350 Holger Way
San Jose, CA 95134 USA
www.force10networks.com

408-571-3500 PHONE
408-571-3550 FACSIMILE

CP24 908 v1.8