

Force10 Networks Delivers Density, Performance to NCSA for 8-Teraflop TeraGrid Cluster

Customer PROFILE

Customer

National Center for Supercomputing Applications

Industry

Scientific Research



Applications

Cluster/Grid Computing

Highlights

A top scientific institution ensures high-performance, cost-effective GbE and 10 GbE Ethernet connectivity at true line rate speeds with E1200 switch/router.



Research institutions such as the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign and the San Diego Supercomputing Center (SDSC) are driving the creation of the next IT revolution: Grid computing. Computational Grids will give users access to enormous "virtual supercomputers" — computers at different locations linked together to work as one — that will provide CPU cycles on demand. One of the most prominent Grid computing projects in the United States is TeraGrid.

Funded by the National Science Foundation, TeraGrid is a multi-year effort to build and deploy the world's largest distributed infrastructure for open scientific research. The TeraGrid will provide the national research community more than 20 teraflops of computing power distributed at the five sites in Illinois, California and Pennsylvania, facilities capable of managing and storing more than 450 terabytes of data, high-resolution visualization environments and toolkits for Grid computing. These components will be tightly integrated and connected through a network that will initially operate at 40 Gbps and later be upgraded to 50 to 80 Gbps — 16 times faster than today's fastest research network.

The capacity of TeraGrid's clusters is only limited by the capacity of the interconnecting network. That's why NCSA — one of four partners in the TeraGrid project — selected switch/routers from Force10 Networks. The Force10 E-Series delivers the high density of Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10 GbE) ports and line rate performance that NCSA needs to maximize the scalability and performance of its network.

High Density in a Single Chassis

With a peak performance of 8 teraflops, NCSA will provide the bulk of the computing power for the TeraGrid using high-powered IBM Linux server clusters running on next-generation Intel Itanium processors. Scientists will be able to use these clusters to run compute-intensive applications in areas such as molecular dynamics, cosmology, weather forecasting and biomolecular electrostatics.

NCSA's plans called for 256 compute nodes in Phase I of the project, which launched in December 2002, all connected by 1000BaseSX Gigabit Ethernet. A key challenge, therefore, was finding a switch/router that could connect this number of systems reliably and still enable the high levels of compute performance — while staying within NCSA's budget.



NCSA 8-Teraflop TeraGrid Cluster

Customer PROFILE

“As far as I’m concerned, the Force10 E1200 has allowed us to create the absolute best Ethernet switching fabric possible within our budget. Our Ethernet infrastructure will enable the high performance operation of our TeraGrid cluster, instead of limiting it. That’s a big win.”

Patrick Dorn
Network Engineer,
National Center for
Supercomputing Applications

"When we were looking at these numbers, a big concern was how many switch/routers it would take to interconnect all of them," says Patrick Dorn, NCSA network engineer. "What really attracted us to Force10 was the density that we could get in a single chassis — and still achieve line-rate performance. It's more cost effective, it's easier for us to manage, and we get better performance."

"Plus, Force10 had the density we needed today," Dorn adds. "With other vendors, it's still on the road map."

Force10 Networks' E-Series supports up to 14 line card slots per chassis, and 24 GbE ports or two 10 GbE ports per line card slot — for a total of 672 ports of GbE or 56 ports of 10 GbE per chassis. For Phase I of its TeraGrid cluster, NCSA installed a Force10 E1200 switch/router fully loaded with four 10 GbE ports and 288 GbE ports. The GbE ports connect the Linux cluster nodes over 1000BaseSX Gigabit Ethernet; the 10 GbE ports connect over 10 GbEBaseLR to the Juniper T640 core router that links NCSA with the TeraGrid backplane.

True Line Rate Performance

Other differentiators that led NCSA to choose Force10 Networks included the E-Series' ability to deliver true line rate performance on all ports, as well as full L2 switching and L3 routing.

"A lot of the ports out there that are advertised as 10 GbE just aren't capable of doing line rate," says Dorn. "They either don't have the backplane connection, or they don't have the forwarding capabilities. In contrast, Force10 has all of those features."

Force10 achieves a high level of performance through several technological innovations in switch fabric, backplane, ASIC design and system control plane. For example, the E1200 E-Series switch fabric provides 56.25 Gigabits per second of non-blocking bandwidth to each line card slot. The high-speed, non-optical backplane is the industry's first to achieve 1.68 Tbps in a single half-rack switch/router chassis. Finally, the Force10 ASICs, along with advanced TCAMs on every line card, give predictable line rate forwarding for every packet regardless of the number, type, or complexity of features enabled across the chassis.

Another important NCSA requirement was jumbo frame support, which allows NCSA to lower interrupt levels as well as the CPU load required to transmit data at faster than 500 Mbps. This means that each node can devote all of its computational cycles to performing computations, rather than transmitting data.

"Force10 took our requirements and worked quickly to meet them," says Dorn. "They gave us a timeframe and didn't slip, which was a phenomenal accomplishment. Their level of responsiveness has been truly impressive."




Figure 1. Force10 enabled NCSA to collapse a 256 node server cluster network into a single E1200 while increasing overall performance and reducing cost.

NCSA 8-Teraflop TeraGrid Cluster

Customer PROFILE

Cost-effective Switching with Room to Grow

To support plans to scale up to nearly 1,000 nodes in Phase II, which is ongoing through 2004, NCSA intends to add another four E1200 systems to its network configuration. The scalability of the Force10 systems gives NCSA the headroom to double port density as the TeraGrid grows. For more information on the NCSA TeraGrid see <http://www.ncsa.uiuc.edu/AboutUs/Hardware/>.

"As far as I'm concerned, the Force10 E1200 has allowed us to create the absolute best Ethernet switching fabric possible within our budget," says Dorn. "Our Ethernet infrastructure will enable the high-performance operation of our TeraGrid cluster, instead of limiting it. That's a big win."



Force10 Networks, Inc.
1440 McCarthy Boulevard
Milpitas, CA 95035 USA
www.force10networks.com

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

© 2004 Force10 Networks, Inc. All rights reserved. Force10, the Force10 logo, EtherScale, FTOS, and TeraScale are trademarks of Force10 Networks, Inc. All other brand and product names are trademarks or registered trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be generally available. Force10 Networks, Inc. assumes no responsibility for any errors that may appear in this document.

CP08 804 v1.5