

San Diego Supercomputer Center Chooses Force10 Switch/Router for TeraGrid Connectivity

Customer PROFILE

Customer
San Diego Supercomputer Center

Industry
Scientific Research



Applications
Cluster/Grid Computing

Highlights
E1200 switch/router delivers the high density of GbE and 10 GbE needed to provide non-blocking, line rate connectivity to Linux clusters.

Enterprises and research institutions once turned to mainframes or supercomputers for the computing power needed to run their most processor- and data-intensive applications. Today, they are turning to clusters and Grids made up of tens to thousands of small commodity servers, interconnected with scalable, high-performance Ethernet networks. These Grids will give users access to enormous "virtual supercomputers" — computers at different locations linked together to work as one — enabling breakthroughs in science and engineering.


One such project in the works is TeraGrid, a multi-year, \$88 million effort funded by the National Science Foundation to build and deploy the world's largest distributed infrastructure for open scientific research. The TeraGrid will bring together distributed scientific instruments, computing facilities with hundreds of thousands of server units, multiple terabyte data archives, and Gigabit networks — all widely accessible by scientists and engineers.

The TeraGrid environment will offer researchers more than 20 teraflops of computing power and nearly 1 petabyte of storage capacity to handle the most complex scientific applications. It will also be extensible and expandable, so that additional sites can connect to it in the future. But the performance of TeraGrid's clusters does have a limit: the capacity of the interconnecting

network. That's why the San Diego Supercomputer Center (SDSC) at the University of California in San Diego — one of five partners in the TeraGrid project — selected switch/routers from Force10 Networks to provide the highest levels of non-blocking connectivity to its Linux server clusters. The Force10 E-Series delivers the high density of Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10 GbE) ports and line-rate performance that SDSC needs to maximize the scalability and performance of its network.

High-density Ethernet Connectivity

The initial phase of the TeraGrid spans five sites in Illinois, California and Pennsylvania: SDSC; the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign; Argonne National Laboratory in Argonne, Illinois; the Center for Advanced Computing Research (CACR) at the California Institute of Technology in Pasadena, California; and the Pittsburgh Supercomputing Center (PSC) at Carnegie Mellon University and University of Pittsburgh, Pennsylvania.



SDSC Chooses Force10 for TeraGrid Connectivity

Customer PROFILE

These facilities will be capable of managing and storing more than 450 terabytes of data, high-resolution visualization environments, and toolkits for grid computing. TeraGrid components will be tightly integrated and connected through a dedicated wide area "backplane" that will initially operate at 40 Gigabits per second and later be upgraded to 50 to 80 Gigabits/second — 16 times faster than today's fastest research network. Clusters that are connected over this grid are being built using 10 GbE interconnections between Gigabit Ethernet-attached Linux server clusters.

SDSC is serving as the lead TeraGrid data and data services site, providing 500 terabytes of SAN disk storage augmented by two 32-processor IBM Power4 DB2 servers joining the existing 72-processor Sun Fire 15K server. The site also will deploy a 4-teraflop Linux cluster of an initial 128 nodes, based on Intel's 64-bit Itanium processor and a 1.1-teraflop Power4 computing system. A major challenge for SDSC, therefore, is building the interconnecting network with the capacity to handle this much computing power.

"It's a density issue," says Nathaniel Mendoza, SDSC lead network test engineer. "We needed a way to connect a large amount of nodes in as small a footprint as possible. Force10 Networks delivers incredibly high density in a single chassis, giving us better scalability and allowing us to install one box instead of several."

"It's a density issue. We needed a way to connect a large amount of nodes in as small a footprint as possible. Force10 Networks delivers incredibly high density in a single chassis, giving us better scalability and allowing us to install one box instead of several."

Nathaniel Mendoza
Lead Network Test Engineer,
San Diego Supercomputer Center

SDSC selected Force10 Networks E1200 switch/routers for their high density of GbE and 10 GbE ports. Based on Force10's E-Series architecture and Force10 ASICs, the Force10 E-Series supports up to 14 line card slots per chassis, and 48 GbE ports or four 10 GbE ports per line card slot — for a total of 672 ports of Gigabit Ethernet or 56 ports of 10 Gigabit Ethernet per chassis. As a result, non-blocking clusters of up to 624 server nodes can be built with a single chassis, which includes 10 GbE uplinks as well. Networked, the E-Series allows non-blocking clusters of well over 3,000 server nodes.

"We chose Force10 because of its unmatched scalability and performance," says Kevin Walsh, senior network engineer at SDSC. "Our requirements are demanding, since we are today aggregating hundreds and in the future thousands of Gigabit Ethernet-attached Linux servers. Force10 is the only vendor we've seen that can meet the challenge and scale with us."

In Phase I of the TeraGrid project, which launched in December 2002, SDSC used one E1200 switch router to provide Gigabit Ethernet connectivity between each node of a 128-node supercomputer cluster. The E1200 then connects over four 10 GbE LR interfaces to the Juniper T640 core router that performs WAN transport between SDSC and the TeraGrid backplane linking all four sites.

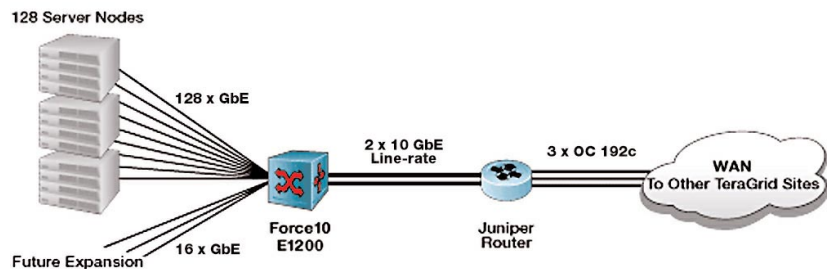


Figure 1. Force10 delivers scalability for future expansion and line-rate performance to SDSC's TeraGrid connected high-performance computing cluster.

SDSC Chooses Force10 for TeraGrid Connectivity

Customer PROFILE

As SDSC's clusters grow in size, the scalability, density and stability of the E1200 will allow the center to scale to very large numbers of nodes. For example, in Phase II, SDSC will connect another 256 nodes to the E1200 and add more 10 GbE LR interfaces — while still maintaining high levels of performance. The center also plans to add another E1200 system in Phase II.

True Line Rate Throughput

Another reason SDSC selected switch/routers from the E1200 series was for their ability to provide true line rate throughput on every port.

"We needed to get as close to wire-rate as possible," says Mendoza. "At the 10 GbE level, most vendors can only deliver 8 Gigabits per slot. Force10 delivers true 10 GbE, reliably, and we still have headroom to double port density in the future."

Force10 achieves this level of performance through several technological innovations in switch fabric, backplane and ASIC design. For example, the E1200 series switch fabric provides 48 Gigabits per second of non-blocking bandwidth to each line card slot, as well as advanced queuing, multicast and jumbo frame support. The high-speed, non-optical backplane is the industry's first to achieve 1.68 Tbps in a single half-rack switch/router chassis. In addition, the Force10 ASICs, along with advanced TCAMs on every line card, deliver predictable line rate forwarding for every packet regardless of the number, type, or complexity of features enabled across the chassis.

A final factor in SDSC's decision to choose Force10 Networks was the company's responsiveness and high level of customer support.

"Force10 gave us unparalleled access to almost anyone in the company," says Mendoza. "They also actively solicited — and incorporated — our input on product direction. They showed an active interest in working with us to deliver a solution that met our needs."



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.

CP09 804 v1.5