

# Berkeley Lab Team Sets Performance Record in 10 Gigabit Ethernet Cluster Computing Test

## Customer PROFILE

### Customer

Lawrence Berkeley National Laboratory

### Industry

Scientific Research



### Applications

Cluster computing demonstration achieves 10.6 Gbps sustained data rate and showcases real-world scientific application for 10 GbE.

### Highlights

LBL used two Force10 E1200s, each aggregating GbE links from 11 dual-CPU servers. One of these clusters runs a simulation of black hole collision while the other runs a visualization application. With line rate performance on all E1200 ports, a sustained 10.6 Gbps communication between two clusters over two 10 GbE links was achieved.

*With the IEEE's recent adoption of standards for 10 Gigabit Ethernet (10 GbE) equipment, high-speed Ethernet has become more than just a theoretical concept — it's moving into the mainstream, and research institutions such as Lawrence Berkeley National Laboratory (LBNL) are leading the way. The laboratory recently proved that 10 GbE can provide the bandwidth today for the most compute-intensive applications, and help organizations cluster computing resources to provide supercomputing power to the desktop.*

Using the E1200 switch/router from Force10 Networks to connect two high-performance Linux clusters, LBNL demonstrated a scientific application that sent terabytes of data from one cluster to another across two 10 GbE connections at combined sustained data rate of 10.6 Gbps.

### Real-world Applications for 10 Gigabit Ethernet

Located in Berkeley, California, and managed by the University of California, Lawrence Berkeley National Laboratory operates some of the world's most powerful computing, data storage and networking resources for the U.S. Department of Energy. The laboratory is a leader in high-performance computing, having won the High-Performance Bandwidth Challenge, held annually at the Supercomputing conference, for two years in a row. In fact, LBNL won the 2001 challenge by moving scientific data across the network at a sustained rate of 3.3 Gbps.

In the summer of 2002, the same team set out to break that record and demonstrate a real-world clustering application for 10 GbE. Mike Bennett, senior network engineer for LBNL's on-site networking services group and part of the team assembling the demonstration, notes that demonstrations of 10 GbE to date have typically used network performance analyzers to generate load. While this shows the ability of network equipment to perform under a limited set of conditions, it hasn't demonstrated the ability of real applications to take advantage of the available bandwidth.

"Many people think that the only useful purpose for 10 Gigabit Ethernet today is to aggregate multiple lower bandwidth LAN segments into a 'fatter pipe,'" says Bennett. "We wanted to show that there are real-world applications that can take advantage of this increased bandwidth now. It's not just theoretical."

## Berkeley Lab



### Clustering Demonstration Surpasses Expectations

The demonstration consisted of two powerful Linux clusters from FineTec Computers, each with SysKonnnect Gigabit Ethernet (GbE) interfaces, connected by a pair of Force10 Networks E1200 switch/routers linked over two 10 GbE interfaces. Ten of the servers connected to one 10 GbE link, while the remaining server connected to the other. Each of the Force10 E1200 systems contained one line card with 48 GbE ports, and one card with two 10 GbE ports. According to Bennett, LBNL selected the E1200 system for its performance, port density, switching capacity and reliability.

# 10 Gigabit Ethernet Cluster Computing Test

## Customer PROFILE

“Force10 equipment works as advertised: it delivers 10 Gigabit Ethernet, reliably and consistently.”

**Mike Bennett**  
Senior Network Engineer,  
Lawrence Berkeley National Laboratory

"In setting up the demonstration, one obstacle was finding network equipment that could deliver at top rates," Bennett says. "But Force10 equipment works as advertised: it delivers 10 Gigabit Ethernet, reliably and consistently. We had used the E1200 system at the Supercomputing 2001 conference, and so we knew that a lot of work had gone into testing the systems. That's why I was comfortable using Force10 systems in the demonstration, because I'd already proven that they had the capacity to deliver."

During the demonstration, one cluster of 11 dual-CPU PCs ran Cactus, a scientific application that simulates the gravity waves resulting from the collision of two black holes. The resulting data was fed over the 10 GbE connections to

the second cluster, where a remote visualization application called Visapult rendered the received data for real-time display and analysis.

The team members not only met their goal of demonstrating sustained 10 GbE performance — they surpassed it. Over a 12-hour period, the two applications used the entire bandwidth of the 10 GbE network, moving 53 terabytes of data over the two 10 GbE links at a combined, sustained data transfer rate of 10.6 Gigabits per second.

"Thanks to Force10, we didn't have to worry about running into some limit and having to engineer around it from a network design standpoint," says Bennett. "The equipment was all top-notch, and Force10 came through for us with excellent support. We proved that the infrastructure exists to support the high-bandwidth connectivity our researchers need today to push the envelope of scientific research."




Figure 1. LBL used two Force10 E1200s, each aggregating GbE links from 11 dual-CPU servers. One of these clusters runs a simulation of black hole collision while the other runs a visualization application. With line rate performance on all E1200 ports, a sustained 10.6 Gbps communication between two clusters over two 10 GbE links was achieved.



**Force10 Networks, Inc.**  
1440 McCarthy Boulevard  
Milpitas, CA 95035 USA  
[www.force10networks.com](http://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.

CP06 804 v1.4