Los Alamos National Lab Smashes Networking Records with Intel's 10-Gigabit Ethernet Server Adapter

SOLUTION SUMMARY

Researchers document world-class performance and throughput, paving the way for fundamental scientific breakthroughs to occur an order of magnitude faster.

CASE HIGHLIGHTS

Profiled Organization: Los Alamos National Laboratory is a leading research lab operated by the University of California for the National Nuclear Security Administration of the U.S. Department of Energy. Research programs focus in areas such as national defense, energy, biomedical science, environmental protection, computational science and materials science.

Challenge: Identify a networking solution that provides consistently fast and reliable performance for high-volume data transfer across multiple computing environments.

Solution: Demonstrate that the Intel® PRO/10GbE LR Server Adapter sets the stage for 10 Gigabit Ethernet to become an all-encompassing technology for local-area, storage-area, system-area, metropolitan-area and wide-area networks (LANs, SANs, MANs and WANs). Researchers analyzed and documented their results in a variety of scenarios using a powerful array of multiprocessor Intel® Architecture servers.

Benefits: Intel's PRO/10GbE LR Server Adapter proved itself as a cost-effective solution for bulk data transfer, delivering record-breaking network performance, high reliability, high efficiency and low end-to-end latency.

The scientific research conducted at Los Alamos National Laboratory is among the most sophisticated in the world. Researchers recently turned their attention to network performance and the newly ratified 10 Gigabit Ethernet (10GbE) standard because examining complex problems in intense detail requires the fastest computing and networking technologies possible.

Intel® PRO Network Connections

Summary

From global climate modeling to HIV genetic sequencing, and from astrophysics analyses to quantum computing, the scientific research conducted at Los Alamos National Laboratory is among the most sophisticated in the world. Scientists at the lab — located on a mountainous 40-square-mile site near Santa Fe, New Mexico — have distinguished themselves by contributing to next-generation research since Los Alamos was created in wartime 1943 to develop the first atomic bombs.

Sixty years later, Los Alamos continues to expand the boundaries of science in dramatic new ways. For example, researchers recently turned their attention to network performance and the newly ratified 10 Gigabit Ethernet (10GbE) standard because examining complex problems in intense detail requires the fastest computing and networking technologies possible. Centering their work on the Intel PRO/10GbE LR Server Adapter, scientists conducted a battery of experiments focused on optimizing 10GbE for networks of workstations, clusters, and grids.

This case study profiles the results of their tests, which culminated in sending a terabyte of data halfway across the world in less than an hour. The achievement — a collaboration with researchers from the California Institute of Technology, CERN, and the Stanford Linear Accelerator Center — was noted in the 2004 Guinness Book of World Records. Los Alamos scientists also compiled their findings in a technical paper **(http://public.lanl.gov/feng/sc03.ps** or **http://public.lanl.gov/feng/sc03.pdf)** that is being presented at the SC2003: High-Performance Networking & Computing Conference in November 2003. The key proof point: Scientists demonstrated that Intel's PRO/10GbE LR Server Adapter delivers superior performance for demanding applications in a full range of network environments.

10GbE FOR THE REST OF US

While scientists at Los Alamos National Laboratory took a hard look at 10GbE from the perspective of raw throughput, savvy IT managers are looking at the technology for compelling performance benefits. For example, moving to the next order of network bandwidth will significantly reduce the time it takes to perform data backup because a 10GbE adapter can move the same document up to eight times faster than a Gigabit Ethernet adapter can.

Organizations today are looking at bandwidth-hungry applications in their environments and upgrading their infrastructure to 10GbE to realize better business productivity. 10GbE is especially attractive in industries that generate and store vast amounts of data, such as finance, publishing and banking, where users perform large amounts of computationally intensive calculations and/or need to transfer large amounts of data back and forth quickly.

Other scenarios likely to see significant benefit from 10GbE bandwidth include high-end enterprise environments and multinational organizations that routinely send information across large-scale networks, as well as businesses that transfer large imaging files. The making of the 2003 animated movie "The Hulk," for example, capitalized on 10GbE to accelerate video-editing processes.

As 10GbE moves into the mainstream, companies can expect two developments: 1) assuming the technology follows the pattern of previous generations of Ethernet, 10GbE should achieve full 10Gbps line rate within the next two to three years, 2) Intel will continue to develop new products that accelerate the transition to higher Ethernet speeds, not only in high-performance scientific applications, but through products that introduce the high-speed benefits of 10GbE connectivity into the mainstream LAN environment.

Scientific Applications Demand High-Performance Networking

In 1975, when early Ethernet technology showed it could deliver 2.94 megabits per second (Mbps) of throughput over a 10 Mbps medium, many were dazzled. But this ubiquitous technology has continued to advance, meeting ever-escalating performance demands. Today, Ethernet arguably carries 99.99 percent of Internet packets, and the performance bar is set at 10 gigabits per second (Gbps), or 10GbE, the standard ratified in 2002. And although 10GbE differs markedly from its first incarnation, the standard continues to ensure interoperability with existing Ethernet, as well as with networking technologies such as SONET, thus setting the stage for Ethernet's expanded used in MANs and WANs.

"Large-scale scientific applications have to move a lot of data between computing nodes, as well as between computing nodes and storage. Typically, this is done via the network, so the more bandwidth you have, the more quickly you can move information and operate on it."

Wu-chun Feng,
Research and Development in Advanced Network Technology, Los Alamos National Laboratory

Scientists at Los Alamos National Laboratory have good reason to want all the networking bandwidth they can get. Whether they're simulating global warming, contributing to genome research, or modeling the composite fracturing of materials, high-speed networking is critical to capturing, transferring, and sharing with colleagues the vast amounts of information they generate.

"Large-scale scientific applications have to move a lot of data between computing nodes, as well as between computing nodes and storage," explains Wu-chun Feng, who leads the Research and Development in Advanced Network Technology (RADIANT) team at Los Alamos. "Typically, this is done via the network, so the more bandwidth you have, the more quickly you can move information and operate on it."

Feng, who led the 10GbE research at Los Alamos, says the difference between 1 Gbps and 10 Gbps of network performance can mean as much as an order of magnitude faster time to solution. Bioinformatics — the science of using computers to solve biological problems — helps put this in real-world perspective. Feng notes that identifying the SARS virus took a global effort, a matter of weeks and enormous volumes of data. Had scientists been empowered to easily share information via distributed "collaboratories" — a computing model that 10GbE enables — the virus might have been identified far more quickly. In the pharmaceutical industry, this could greatly accelerate the drug-discovery process.

To analyze 10GbE performance over networks of workstations, clusters and grids, researchers at Los Alamos designed a series of experiments focusing on bulk data transfer and then used industry-standard tools to measure the results. To facilitate their tests, scientists chose the new Intel PRO/10GbE LR Server Adapter, which supports long-range data transport over distances of up to 10 kilometers using single-mode fiber optic cabling. Feng explains that Intel's PRO/10GbE LR Server Adapter enables high-performance networking over standard Ethernet-based TCP/IP networks.

"Los Alamos could use the Intel PRO/10GbE LR Server Adapter as the basis for making 10GbE an all-encompassing technology from LANs and SANs to MANs and WANs," Feng says. "Intel's PRO/10GbE LR Server Adapter provides a general-purpose, TCP/IP-based solution to applications — a solution that requires no modification to application codes to achieve high performance."

Record-Breaking Performance, Cost-Effective Solution

The Intel PRO/10GbE LR Server Adapter passed the exhaustive tests designed by Los Alamos researchers with flying colors. With appropriate optimizations to the configurations of the 10GbE adapter and TCP, scientists determined that Intel's PRO/10GbE LR Server Adapter performs exceptionally well in LANs, SANs and WANs.

"For local-area, storage-area, and system-area networks in support of networks of workstations, network-attached storage and clusters, respectively, we achieved end-to-end throughput approaching 7 Gbps between applications running on Linux-based PCs, with end-to-end latency as low as

12 microseconds," Feng says. He notes that low latency is especially important in scientific applications such as global climate modeling and hydrodynamics, but it is also important for enterprise applications such as server clustering and storage-area networks.

"The performance offered by the Intel PRO/10GbE LR Server Adapter enabled us to smash the Internet2 Land Speed Record[1] by sustaining end-to-end TCP/IP throughput of 2.38 Gbps across a distance of 10,037 kilometers."

Wu-chun Feng,
Research and Development in Advanced Network Technology, Los Alamos National Laboratory

Researchers evaluated the performance of the Intel PRO/10GbE LR Server Adapter in three LAN environments: a direct single flow between two PCs connected back-to-back via a crossover cable; an indirect single flow between two PCs through a switch; and multiple flows through a switch. The adapters were hosted in Dell PowerEdge* 2650 and 4600 servers. Each PowerEdge 2650 system contained dual 2.2 GHz Intel® Xeon™ processors running on a 400 MHz front-side bus, and each PowerEdge 4600 system contained dual 2.4 GHz Intel Xeon processors running on a 400 MHz front-side bus. With these PCs, researchers achieved 4 to 5 Gbps of throughput and 20-microsecond end-to-end latency; with faster computers based on Intel® Itanium® 2 processors, the numbers improved to 7 Gbps throughput and 12-microsecond latency.

Los Alamos researchers, in conjunction with the California Institute of Technology, CERN, and the Stanford Linear Accelerator Center, also analyzed the performance of Intel's 10GbE adapters across a WAN — and landed in the Guinness World Book of Records. The achievement was noted electronically on the Guinness World Records Web site, **http://www.guinnessworldrecords.com/index.asp?id=58445**, on July 7, 2003.

"The performance offered by the Intel PRO/10GbE LR Server Adapter enabled us to smash the Internet2 Land Speed Record[1] by sustaining end-to-end TCP/IP throughput of 2.38 Gbps across a distance of 10,037 kilometers (approx. 6,236 miles) between Sunnyvale, California, and Geneva, Switzerland," Feng says, noting that this translates to moving more than a terabyte of data roughly halfway across the globe in less than an hour.

In traversing that distance, the WAN traffic crossed two systems — a 10 Gbps circuit from Sunnyvale to Chicago and a 2.49 Gbps transatlantic circuit between Chicago and Geneva. At each end point, scientists stationed a dual 2.4 GHz Intel Xeon processor-based PC with 2 GB of memory and a dedicated 133 MHz PCI-X bus for the Intel PRO/10GbE LR Server Adapter (see illustration).

The WAN results are remarkable for several reasons. First, they demolished both the single- and multi-stream Internet2 Land Speed Records by 2.5 times. Second, the end-to-end WAN throughput was actually greater than what an application user typically sees in a LAN/SAN environment. Third, given the fact that the bandwidth bottleneck was the 2.49 Gbps transatlantic

WAN Test

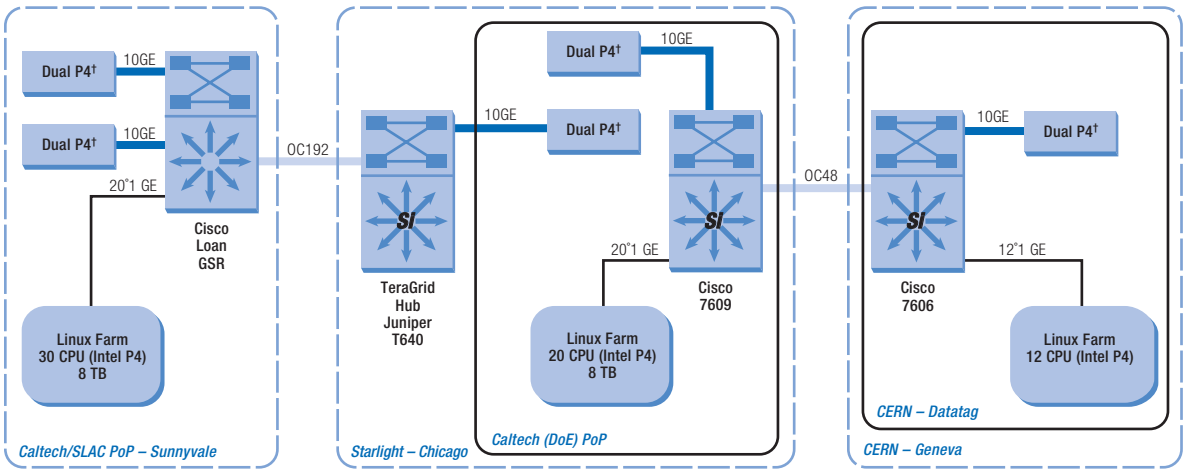

† **Dual 2.4 GHz or 2.2 GHz Intel® Pentium® 4 Xeon™ processors**

circuit, achieving 2.38 Gbps throughput meant that the connection operated at roughly 99 percent efficiency (for the data payload).

"Our combined results indicate that the Intel PRO/10GbE LR Server Adapter is a cost-effective solution for a multitude of computing environments," Feng says.

10GbE can span multiple environments, unlike "exotic" and expensive interconnects that require administrators to maintain two infrastructures – the "exotic" source-routed SAN and the TCP/IP-based network around the SAN. Whether 10GbE is deployed on LANs, supercomputing clusters or a computational grid that spans half the globe, nothing changes. That means applications don't have to be rewritten; they all use the same network interface and the same underlying technology.

10GbE Bodes Well for Future Bandwidth Needs

Scientists recognize that bandwidth needs will only continue to rise. Consider that physicists today must move enormous amounts of data, measured in petabytes,[2] so they can collaborate on experiments with their colleagues around the globe. To comprehend how much disk space a petabyte requires, Feng offers an analogy using time: 1 gigasecond is 31.7 years, while 1 petasecond equals 31.7 *million* years. The amount of bandwidth required to move a petabyte of data over a network would seem equally staggering, but Feng notes that 10GbE technology enables it today.

Scientists, however, aren't the only users clamoring for bandwidth. For example, says Feng, the same technology that displays frames for a bandwidth-driven global climate modeling simulation also applies to the video frames in multimedia streaming, an application that's highly appealing to the general public. Consider that it's already possible to stream an entire feature-length movie halfway across the world in less than 20 seconds. And with enough bandwidth — namely 10GbE — the recipient could download roughly the first 10 minutes of the movie in the first second, and start viewing the movie right away while the rest of the video finished streaming.

The Los Alamos 10GbE research also has implications for other aspects of the PC arena. Feng cites examples such as users who want to play distributed video games or share music with friends around the world.

"Much of the work we do in high-performance networking is not too far ahead of what the general public will see in five years," Feng says. "In other words, what's high performance today becomes commodity tomorrow."

That means users can probably bank on 10GbE following a trajectory similar to that of Gigabit Ethernet, Feng adds. In other words, as product volumes increase when the technology goes mainstream, 10GbE could become a must for exceptional home and enterprise computing environments.

"Los Alamos could use the Intel PRO/10GbE LR Server Adapter as the basis for making 10GbE an all-encompassing technology from LANs and SANs to MANs and WANs, Intel's PRO/10GbE LR Server Adapter provides a general-purpose, TCP/IP-based solution to applications — a solution that requires no modification to application codes to achieve high performance."

Wu-chun Feng,
Research and Development in Advanced Network Technology,
Los Alamos National Laboratory

For more information visit www.intel.com/network

1. Internet2 Land Speed Records are maintained by the Internet2 consortium (http://www.internet2.edu). On October 1, 2003 Caltech and CERN beat the record transferring over a Terabyte of data in less than 30 minutes between Geneva and Chicago. This record was also set using the Intel® PRO/10GbE LR Server Adapter and the sender used dual Intel® Itanium® 2 processors and the receiver used dual 3.06 GHz Intel® Xeon® processors. For more information visit http://lsr.internet2.edu/history.html.

2. "Peta" describes the next thousand-fold order of magnitude after "tera" (which is 10^{12}). A petabyte is therefore 10^{15} or 1,000,000,000,000,000 bytes of data.

Printed in U.S.A. 1103/OC/RC/PP/PDF Please Recycle. 254317-001