




By Jordan Plawner
Travis Vigil

ACCELERATING APPLICATION TRAFFIC WITH INTEL 10 GIGABIT ETHERNET SERVER ADAPTERS

By using Intel® 10 Gigabit Ethernet server adapters in Dell™ PowerEdge™ servers and unified storage systems such as the Dell PowerVault™ NX1950, organizations can achieve the performance required for high-bandwidth applications and meet the demand for increased storage capacity while moving toward network convergence on a simplified Ethernet infrastructure.



For More Information:

[Click here](#) to view the article on *Configuring iSCSI Remote Boot on Dell PowerEdge™ Servers with Intel® Adapters*.

Related Categories:

Dell PowerEdge servers
Dell PowerVault storage
Ethernet
Intel
Intel networking
Internet SCSI (iSCSI)
Network interface card (NIC)

Visit DELL.COM/PowerSolutions for the complete category index.

The growing demand for high-bandwidth, real-time applications by enterprises and consumers can require IT departments to make significant changes in their IT infrastructures. The increasing use of radio-frequency identification (RFID) technologies, decision support software, data mining, virtualized applications, and multimedia content is already pushing internal networks to their limits. At the same time, the rising adoption of Internet software distribution, high-volume data transactions, software as a service, and communications applications is creating additional infrastructure burdens. These trends are likely to accelerate. Internet communications applications, for example, are being used with increasing frequency among both home and business users: a recent survey found that three out of four U.S. Internet users streamed video online in May 2007.¹ Organizations need to find new ways to deliver more data to more people in more locations than they can right now.

Supporting high-bandwidth, real-time applications requires IT groups to add not only bandwidth to their infrastructure, but also storage capacity. Organizations are already facing the challenges of scaling storage capacity to accommodate the growing volume of health care records, financial data, and customer information

required for data mining, customer relations, and e-commerce. The creation, distribution, and archiving of high-resolution content, including the media that people are increasingly streaming from the Internet, will likely require enterprises to scale capacity further. One market research firm recently predicted an eightfold increase in new storage capacity shipments per year between 2005 and 2012.²

As organizations begin to modify their infrastructures to accommodate these trends, they will likely need to do so without significantly adding to the cost or complexity of IT. Enterprises will also likely need to increase storage capacity without major expenditures in hardware acquisition or management, and to prepare their networks for increased throughput without adding complexity through the management of multiple networks.

SUPPORTING HIGH-BANDWIDTH, REAL-TIME APPLICATIONS

By combining the capabilities of a network attached storage (NAS) system and Internet SCSI (iSCSI), unified storage systems such as the Dell PowerVault NX1950 can provide organizations with cost-effective ways to scale storage capacity and address the challenges of supporting high-bandwidth, real-time

¹ "3 Out of 4 U.S. Internet Users Streamed Video Online in May," by comScore, July 17, 2007, www.comscore.com/press/release.asp?press=1529.

² "2007 Digital Storage in Entertainment Creation and Distribution," by Coughlin Associates, available at www.tomcoughlin.com/techpapers.htm.

applications. Unified systems can help deliver the performance required for real-time applications while also helping simplify storage administration and reduce the cost and complexity of adding storage capacity.

Unified storage systems can help optimize information processing for real-time applications by removing processing burdens from the host server. The PowerVault NX1950, for example, stores both file and application data in a single system. While client systems can access file data using Common Internet File System (CIFS) and Network File System (NFS) standards, servers can access application data through the iSCSI protocol. With iSCSI technology, blocks of information are transmitted over an IP network between the server and storage system. Application workloads are managed by host servers, and storage workloads are handled by the PowerVault NX1950. This division of labor enables servers to devote more resources to real-time applications than they can with other types of storage.

A unified storage system can also help simplify storage administration and reduce costs. While Fibre Channel storage systems such as Dell/EMC CX series storage area networks (SANs) offer the performance, reliability, and functionality required by many enterprise functions, the cost of acquiring and maintaining Fibre Channel storage places these systems beyond the reach of some organizations. A unified storage system that employs iSCSI can offer a cost-effective, simplified alternative to a Fibre Channel system. Administrators of iSCSI systems can use standard, cost-effective Ethernet switches, routers, cables, and network interface cards (NICs). Consequently, the overall cost of acquiring and maintaining an iSCSI system can be less than that of a Fibre Channel system with comparable capacity. Because many organizations already manage an Ethernet network, moving from a combination of Fibre Channel and Ethernet networks to a single Ethernet network can also reduce the complexity of network administration.

At the same time, the use of iSCSI storage can bring added advantages over Fibre Channel systems. For example, the use of Ethernet and TCP/IP technology by iSCSI storage can help extend a system's geographic flexibility beyond the limits of a Fibre Channel system. Consequently, organizations can incorporate iSCSI storage into storage centralization and disaster recovery strategies. iSCSI systems can also help organizations realize the benefits of using shared storage with virtualization without incurring the cost or complexity of deploying a Fibre Channel SAN.

In addition, a unified storage system that uses iSCSI can provide administrators with a cost-effective way to capitalize on boot capabilities offered by SANs. By maintaining OS images on a shared storage system, administrators can streamline server provisioning, easily deploy upgrades and patches, and accelerate disaster recovery. A unified system such as the PowerVault NX1950 can provide these advantages without the infrastructure costs of a Fibre Channel-based SAN.

BUILDING TOMORROW'S INFRASTRUCTURE

While many enterprises find that the Gigabit Ethernet bandwidth used by typical unified systems is sufficient for running current applications, other organizations

are adopting servers with Gigabit Ethernet LAN on Motherboard chips in conjunction with multi-port adapter cards to take advantage of multi-gigabit speeds. Some of today's streaming media, data backup, and decision support applications require that increased throughput.

The introduction of Intel 10 Gigabit Ethernet (10GbE) server adapters enables organizations to use unified storage to meet the needs of today's throughput-intensive applications while helping build a foundation for future growth. By using a variety of technologies to distribute packets across multiple processor cores, these adapters provide more than additional bandwidth. They also help accelerate processing for real-time applications.

OPTIMIZING PERFORMANCE WITH MULTI-CORE PROCESSORS

To provide the level of performance required for high-bandwidth, real-time applications, an infrastructure needs to deliver not only the necessary bandwidth capacity but also the means to accelerate the processing of network data packets. A network packet undergoes considerable processing between the time it arrives at the physical controller to the moment its payload is delivered to the application. On many of today's systems, this processing is performed by a single


INTEL ADAPTER TECHNICAL SPECIFICATIONS

Intel 10 Gigabit Ethernet (10GbE) adapters provide comprehensive support for multiple IEEE standards. Three types of fiber-optic adapter are currently available:

- Single-port SR (short reach: up to 300 m)
- Dual-port SR
- Single-port LR (long reach: up to 10 km)

Key technical specifications of Intel 10GbE adapters include the following:

- 32 transmit queues and 64 receive queues per port
- 16 virtual machine device queues per port
- 512 KB receive buffer (divisible into eight individual packet buffers)
- 320 KB transmit buffer (also divisible into eight individual packet buffers)



processor. But multiple I/O requests to a single processor can create a bottleneck while leaving its other resources available for I/O processing.

This problem cannot be solved simply by adding bandwidth. In fact, adding network controllers only amplifies the problem: additional packets increase both processor congestion and latency. The solution lies in the ability to distribute data packet processing across multiple cores working in parallel. By capitalizing on the processing capabilities of these cores, Intel server adapters distribute the load and help lower the processing time for individual network packets.

Intel 10GbE adapters are specifically designed for multi-core systems and the latest OS features. They can use up to 32 separate hardware NIC queues for transmission and up to 64 queues for receiving—all of which can be mapped to 16 separate processor cores.

The ability to direct streams to different processor cores can provide important advantages in virtualized server environments. By using Intel 10GbE adapters, virtual machines (VMs) hosted by hypervisors that emulate network controllers can rely on a dedicated network stream handled by a single processor core. When multiple VMs are in use, they can share the controller ports while taking advantage of their own privately processed packet stream. This approach enables significant enhancements in the performance of virtualized applications.

As illustrated in Figure 1, Intel 10GbE adapters use multiple technologies that work together to distribute packets across multiple processor cores and help accelerate processing:

- **Multiple descriptor queues:** To spread the workload across multiple processor cores, network traffic streams are divided into queues through receive-side scaling (RSS), filtering based on the Media Access Control (MAC) address, or the use of virtual LAN tags. The packet queues can be accessed by

driver threads running on different processor cores so that multiple cores can process network packets in parallel. The Intel 10GbE adapters' 32 transmit queues and 64 receive queues can be mapped to up to 16 processor cores. For servers equipped with several multi-core processors, multiple descriptor queues can facilitate powerful load-balancing functionality.

- **Receive-side scaling:** To determine which queue to use for incoming packets, network adapters residing on systems using the Microsoft® Windows Server® 2003 or Windows Vista® operating systems can use RSS, which directs packets to different queues without the need for reordering. (On systems using the Linux® OS, this technology is known as scalable I/O). RSS is intelligent in its distribution of packet processing, and is also programmable. Consequently, network controllers with multiple queues can efficiently direct multiple TCP/IP streams to different processor cores for handling.
- **Message-Signaled Interrupts Extended (MSI-X):** As part of the PCI Express (PCIe) standard, MSI-X technology facilitates efficient communication between queues and specific processor cores, enabling the network adapter

hardware to direct an interrupt to the designated core when network packets are placed in the processor queue. In this way, MSI-X provides an important enhancement to standard MSI, which could pass interrupts only to a single processor core. Intel 10GbE adapters give each queue its own set of MSI-X controllable interrupt vectors to help provide efficient packet management and fine-tune the processor load.

- **Intel Virtual Machine Device Queues (VMDq):** VMDq provides multiple hardware queues and offload features to help reduce the software overhead associated with sharing a single network controller between multiple VMs. Previously, a network switch emulated by the virtualization software sorted and routed packets individually to VMs. That process typically introduced significant delays in the network packet processing. With VMDq, individual hardware queues are associated with the simulated network interfaces of running VMs. The network controller itself performs the routing of received packets, helping substantially reduce overhead. VMDq is also used on outbound VM packets to provide transmit fairness and to help avoid a single VM blocking access to the controller.

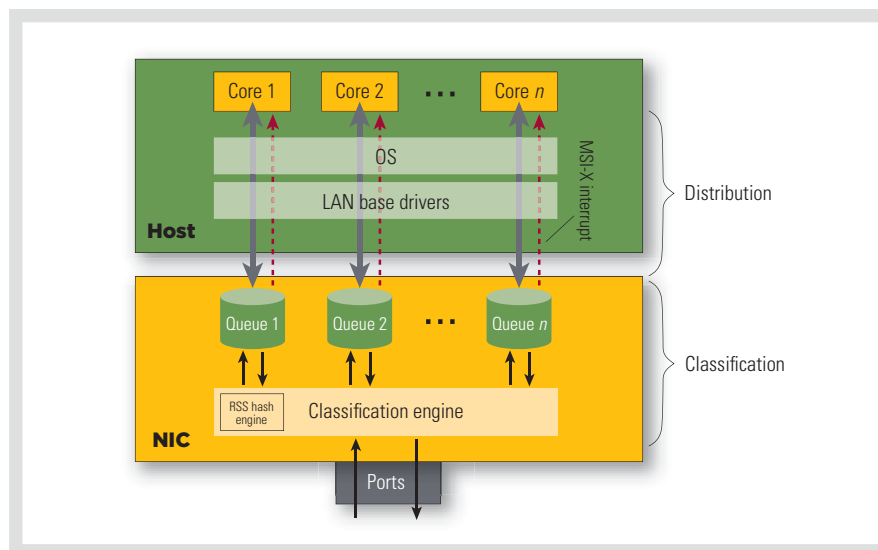


Figure 1. Intel 10 Gigabit Ethernet adapters sort, group, and direct network packets across multiple processor cores to help reduce network bottlenecks




Figure 2. Dell PowerEdge 2950 servers connected to Dell PowerVault NX1950 storage using Intel 10 Gigabit Ethernet adapters help deliver a cost-effective file-sharing solution that can support high-bandwidth enterprise applications

■ Intel I/O Acceleration Technology:

Intel 10GbE adapters also include Intel I/O Acceleration Technology, which helps optimize bandwidth by redirecting header processing and speeding up memory access to packet components.

For more information on the technical aspects of the Intel 10GbE adapter hardware, see the “Intel adapter technical specifications” sidebar in this article.

EXPANDING THE DELL 10 GIGABIT ETHERNET ECOSYSTEM

Intel 10GbE adapters are validated for the Dell PowerVault NX1950 storage system and for Dell PowerEdge servers, while Dell PowerConnect™ 6200 series switches provide the connectivity for a Dell 10GbE infrastructure.

Dell PowerVault NX1950 storage system

The Dell PowerVault NX1950 network storage system offers a range of advantages for organizations that need a cost-effective way to address the growing demand for storage capacity and to support high-bandwidth, real-time applications (see Figure 2). While the unified design helps optimize the processing power of host servers, the simple, cost-effective scalability of the PowerVault NX1950 can help address rapidly rising demand for storage capacity. The

integrated PowerVault MD3000 modular disk storage array holds up to 15 Serial Attached SCSI (SAS) drives for up to 4.5 TB of storage capacity. Administrators can also easily add up to two PowerVault MD1000 external storage expansion arrays, for a total of 45 drives and up to 13.5 TB of total storage capacity in a single 3U enclosure. In many cases, these SAS drives are less expensive than comparable Fibre Channel disks.

The PowerVault NX1950 also helps simplify storage deployment and management. While setup wizards help administrators configure the system quickly, the included Microsoft Windows® Unified Data Storage Server 2003 OS provides an integrated console in a familiar management environment for administrators to conduct essential tasks.


Dell PowerEdge servers

Intel 10GbE adapters are validated and available as an option on ninth-generation Dell PowerEdge servers. By offering a selection of multi-core processors, including dual- and quad-core Intel Xeon® processors on several models, PowerEdge servers can provide the performance required for real-time processes. The Intel adapters add cost-effective iSCSI connectivity while helping optimize the network processing performance of multi-core processors to deliver the performance required for high-bandwidth applications.

Dell PowerConnect 6200 series switches


Advanced Dell PowerConnect 6200 series switches can incorporate 24 or 48 ports of 10/100/1,000 Mbps Ethernet in a 1U form factor, enabling administrators to connect up to 576 servers or clients and up to nearly 3 TB of capacity in a single stack. These switches also support up to four 10GbE uplinks for direct connectivity to 10GbE routers, enterprise backbones, and data centers. The modular design enables administrators to upgrade to advanced stacking or 10GbE when needed.

PREPARING FOR THE FUTURE

Unified storage systems such as the Dell PowerVault NX1950 offer organizations a cost-effective, scalable, easy-to-manage solution for the growing demand for storage capacity, while Intel 10GbE server adapters help increase the viability of iSCSI storage even further for enterprise data centers. With increased bandwidth and technologies designed to accelerate network processing by capitalizing on the capabilities of multi-core processors, Intel 10GbE server adapters can help prepare organizations for today's and tomorrow's high-bandwidth, real-time application requirements. 

Jordan Plawner is a product planner and technologist in the Intel LAN Access Division.

Travis Vigil is a product marketing strategist for Dell iSCSI and Dell PowerConnect solutions.



ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Intel 10GbE server adapters:
www.intel.com/go/ethernet

Dell PowerEdge servers:
DELL.COM/PowerEdge