Maximizing Servers Through Multiple Gigabit Connections

By Gary Gumanow

As enterprise networks upgrade to powerful desktop systems that can simultaneously move large files and handle multiple sophisticated applications, existing network server connections may not be able to meet the increased bandwidth demands. This article details several methods for enhancing server performance and availability, including the use of dual Gigabit LAN on Motherboard (LOM) technology available in select Dell® PowerEdge® servers.

Sophisticated applications and increasingly powerful processors are driving network traffic to new levels, compelling re-evaluation of critical connections. Robust, server-centric applications and the proliferation of Gigabit Ethernet (GbE) to the desktop can hasten a need for multi-Gigabit network connections for enterprise servers. Without these connections, users experience slow response times, sluggish retrieval of files stored on the server, and significant loss of productivity.

Now nearing standards approval by the IEEE (Institute of Electrical and Electronics Engineers), 10 GbE should provide more than enough bandwidth for applications and servers supporting 1 GbE connections to the desktop. However, server adapters based on the IEEE® 802.3ae standard are not available today.

In the meantime, network managers can create multi-Gigabit server connections without 10 GbE adapters. Methods for creating these connections include network segmentation, fault tolerance, and adapter teaming/load balancing. Together, these methods represent an escalating sequence of actions that network managers can consider as networking demands increase.

Segmenting a network and building fault tolerance

Network managers have often solved server or network bottlenecks by installing additional adapters in the server and segmenting the network into multiple subnets, each with its own dedicated adapter (see Figure 1). Although this method requires additional hardware and management overhead including repeated reconfiguration to balance the traffic load, it is effective in reducing the traffic volume on each network link. Many network administrators currently use this method.

One drawback to segmenting a network is that it forces the enterprise server to become responsible for distributing packets

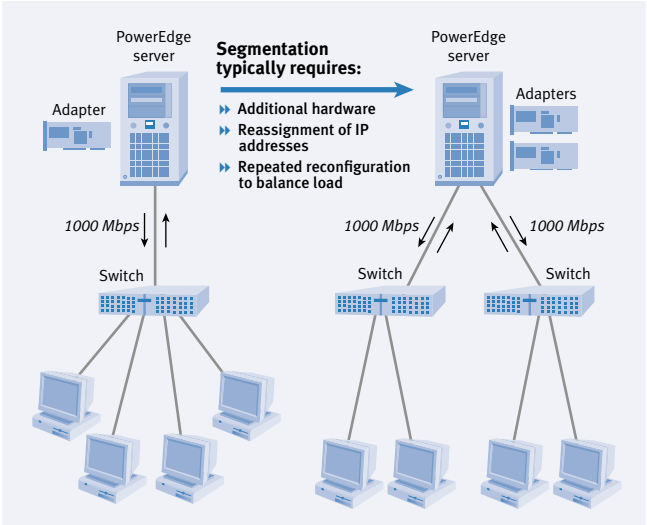

Figure 1. Segmenting a network to increase server bandwidth

Advanced features	Connection requirements	Benefits
Adaptive Load Balancing	When connected to any switch	• High server availability • Load balancing of outgoing server traffic for increased throughput up to 800 Mbps (8 Gbps in Gigabit Ethernet environments)
Intel Link Aggregation, Fast EtherChannel (FEC), or Gigabit EtherChannel (GEC)	When connected to a switch that supports Intel Link Aggregation, FEC, or GEC	• High server availability • Load balancing of outgoing and incoming server traffic, for throughput up to 1600 Mbps (potentially 16 Gbps with Gigabit Ethernet)
IEEE 802.3ad	When connected to a switch that supports IEEE 802.3ad link aggregation	• Higher server availability • Balances traffic among multiple switches from one server for higher throughput and increased redundancy

Figure 2. By configuring the adapter for the desired advanced feature support, companies can match the needs and resources of their particular server environment

to the individual adapters and their respective subnets. Furthermore, this architecture provides no redundancy if a link fails. Administrators implementing network segmentation can maximize fault tolerance and availability by ensuring that multiple adapters in a server are configured for automatic failover from one adapter to another.

Because today's organizations rely more and more on powerful client systems and demanding applications, they should consider configuring Gigabit server adapters to fail over to GbE, not 10/100 Mbps Ethernet, to avoid network performance degradation. To accommodate the proliferation of GbE desktops, servers such as the Dell® PowerEdge® are now available with dual GbE LAN on Motherboard (LOM) connections for built-in redundancy without loss of performance.

Delivering the headroom for multi-Gigabit

Growing bandwidth needs will soon require more than two Gigabit Ethernet connections. An easy, cost-effective implementation that addresses this need is to add a single-port or dual-port GbE adapter to the GbE connections already built into a dual-LOM server. The result is multi-Gigabit performance while using only a single peripheral component interconnect (PCI) or PCI Extended (PCI-X) slot. This method leaves valuable PCI slots available for other purposes ranging from RAID (redundant array of independent disks) disks to Secure Sockets Layer (SSL) encryption cards.

Intel® Labs used the Large Send Offload (LSO) feature of the Intel PRO/1000 Gigabit Server Adapters to test three 1 GbE adapters working together. This test scenario, which also used Microsoft® Windows® 2000 Advanced Server and Microsoft Windows XP™ operating systems, achieved 3 Gbps throughput. By off-loading and accelerating the packet formation process, LSO reduces the host CPU burden, which leaves available headroom for delivering multi-Gigabit.

For the most demanding situations, the ideal way to move more data between the server and the network is to install a team of GbE adapters and automatically balance traffic across them. Each adapter adds more bandwidth, and the automatic distribution of traffic across the adapters eliminates the need to segment the network.

Creating one logical pipe

Several proven technologies are available to provide scalable server bandwidth through load balancing as well as to provide automatic, redundant connections for increased server availability (see Figure 2). These technologies treat the adapter team as one logical pipe.

An example of such technology is the IEEE 802.3ad industry standard for link aggregation, a standard that Intel supports. It provides compatibility with other vendor switches that support the standard and allows for balancing traffic among multiple server adapters and switches (see Figure 3).

An easy, cost-effective implementation is to add GbE adapters to the GbE connections already built into a dual-LOM server.

Adaptive Load Balancing (ALB), a technology developed by Intel, also can increase server bandwidth by automatically balancing traffic across as many as eight server adapters. Each additional GbE adapter adds another 1000 Mbps link to the network. Since the distribution of traffic among the adapters is automatic, network segmentation or reconfiguration is unnecessary. All the adapters share the existing IP address and ALB continuously balances traffic among the adapters.

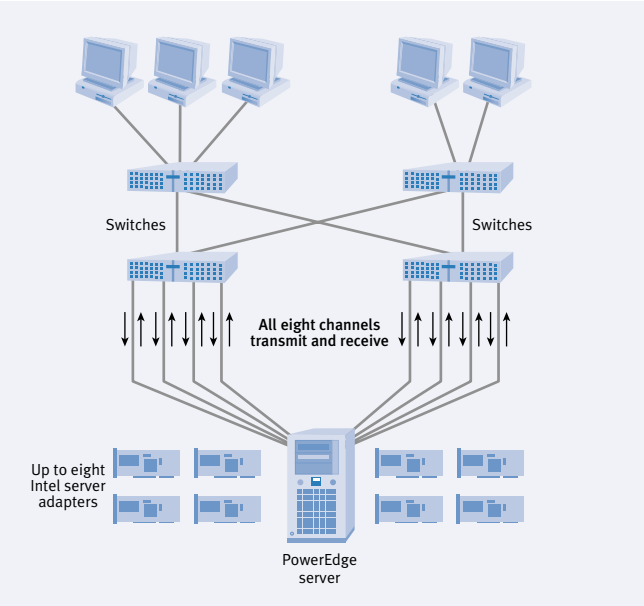

Figure 3. Network traffic with IEEE 802.3ad link aggregation

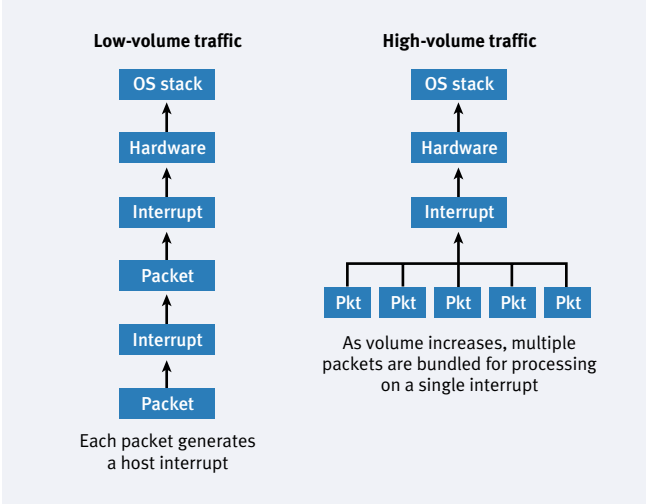

Figure 4. Packet bundling for host interrupt moderation

Another technology, Gigabit EtherChannel® (GEC), developed by Cisco Systems and supported by Intel, supports scalable bandwidth for up to eight adapters at full duplex and potentially can deliver in excess of 8 Gbps. This technology requires support in both the adapter and the connected switch.

Exploring the present and future of interrupt moderation

Future Intel Gigabit server adapters and LOM will provide advanced interrupt moderation technology, which will additionally reduce host processor interrupts. This reduction in interrupts enables Gigabit EtherChannel technology to deliver even more of its 16 Gbps bandwidth potential (8 Gbps × full duplex). The load-balancing application required for teaming, Advanced Network Services, will have much less impact on overall performance because the server CPU will have more headroom.

To achieve optimum throughput, it is necessary to control host processor interrupts, which the adapter generates to request cycles for packet processing. Too few interrupts can lead to latencies and too many can unduly burden the server CPU. Bundling an appropriate number of packets before issuing an interrupt to the host (see Figure 4) allows the Intel Gigabit server adapter to "tune" the interrupt frequency to match traffic conditions while still maintaining packet flow.

As this and similar technologies become available, network managers should explore interrupt moderation methods that use dynamic timing adjustment, not just sampling; the more granular the timing, the better the performance.

For most networks, the technologies discussed in this article can provide the multi-Gigabit bandwidth needed for server and backbone connections and can support movement to Gigabit Ethernet at the desktop. Even after 10 GbE server adapters become available, link aggregation, load balancing, and interrupt moderation technologies should continue to provide value. At the very least, they can enable organizations to implement a smoother and more cost-effective migration to 10 GbE networking speeds.

Even after 10 GbE server adapters become available, link aggregation, load balancing, and interrupt moderation technologies should continue to provide value.

Gary Gumanow *(gary.gumanow@intel.com) is a senior technical marketing engineer in the Platform Networking Group at Intel Corporation. He is responsible for product planning for Intel Gigabit Ethernet adapters and software. Gary has been designing networks for more than 18 years and holds an MBA from Pace University.*

FOR MORE INFORMATION

For more information, visit http://www.intel.com