



Using Segmentation to Increase Network Performance

More network users, greater application complexity, larger file sizes—all of these can increase network traffic, resulting in decreased network performance. Using multiple Gigabit Ethernet server ports provides a quick and cost-effective method for segmenting traffic loads to enhance network bandwidth, response and reliability.

December 2004

Contents

Executive Summary	2
Addressing the Burdens of Enterprise Growth	3
Boosting Performance with Segmentation	3
Switches Enhance Control and Reliability	5
Combining NICs and LOM for Performance	6
Further Performance Enhancements	7
Conclusion	8

Executive Summary

Business growth, while generally desirable, imposes additional burdens that can diminish network performance. Performance-robbing burdens include rapid growth in the number of network users or clients, higher traffic volume generated by each client, increasing application complexity and larger file sizes.

This white paper discusses network segmentation as a quick and cost-effective means of boosting network performance in any business, from rapidly growing start-ups to expanding multinational corporations. The role of Intel® Server Adapters is also discussed, both as a tool for implementing segmentation and as a means for enhancing reliability and server efficiency. Other topics include the use of multi-port adapters to conserve server slots and the benefits of on-board adapter processing to free server resources.

Addressing the Burdens of Enterprise Growth

Ideally, network growth is well planned and orchestrated in advance to meet the additional burdens of business growth. Too often, though, network expansion is reactive rather than proactive. Business growth just isn't that predictable. This is especially true for rapid growth in small start-ups, for consolidating branch offices into a regional office, or even in the acquisition-fueled growth of a multinational corporation—size is no antidote for growth pains. While information technology (IT) managers try to plan for network growth, rapid or unexpected growth can stretch IT budgets and staff to the limit. Then the overriding objective becomes, "Let's just get something up and running right now." This, in turn, may result in missing the full potential for realizing greater network performance and scalability, leaving the IT staff to face major performance enhancements in the future.

Fortunately, IT managers can mitigate the burdens of rapid growth by using network segmentation. Segmentation is a quick and cost-effective method for multiplying network bandwidth, depending on the number of segments used. Additionally, segmentation's performance gains can be further augmented by various other upgrades, including upgrading servers from Fast Ethernet (100 Mbps) to Gigabit Ethernet (GbE), replacing dumb hubs with smart switches, and adding switched redundant links for higher network reliability. For high-performance demands, especially for application-intensive operations with numerous transfers of extremely large files, network administrators may want to consider converting their network backbone to 10 Gigabit Ethernet (10GbE).

The first step in performance enhancement, however, is typically network segmentation. This is particularly true for small start-ups or branch offices. In such cases, networks tend to start out small and simple—a single server with ten or fewer clients, for example—where segmentation hasn't been implemented. As the organization grows or applications become more complex, segmentation becomes a necessity for improving network response, throughput and resiliency. Additionally, segmentation opens growth paths for network organization according to workgroup needs and for network expansion to accommodate multiple workgroups in large buildings.

Boosting Performance with Segmentation

Segmentation's basic function is to split traffic loads, thus alleviating bottlenecks. This is, in essence, comparable to changing a two-lane highway into a four-lane highway. More traffic flows quicker. Add another segment—more lanes—and even more traffic can flow quickly.

In addition to enhancing throughput, segmentation offers network administrators other advantages, such as high security and reliability, which can be gained by careful definition of segments and judicious selection of the implementing hardware. To gain insight into these additional advantages and the throughput gains offered by segmentation, let's take a brief look at a simple network topology and evolve it into a campus network.

First, consider the basic network shown in Figure 1, which is typical of an embryonic enterprise. This is the seed from which even the largest enterprise networks once germinated. However, Figure 1 shows a key difference. Notice that the simple network of Figure 1 has a Gigabit Ethernet (GbE) trunk. Many servers, especially in older networks, still use a Fast Ethernet adapter. As a first performance upgrade, such servers should be converted to Gigabit Ethernet. This conversion can be done quickly and easily from a wide choice of Intel® PRO/1000 Server Adapters for both copper cabling (CAT 5) and fiber-optic cabling. Converting servers from Fast Ethernet (100 Mbps) to Gigabit Ethernet provides an immediate bandwidth advantage for higher throughput.

Figure 1. Basic embryonic network topology





Figure 2. Basic network segmentation

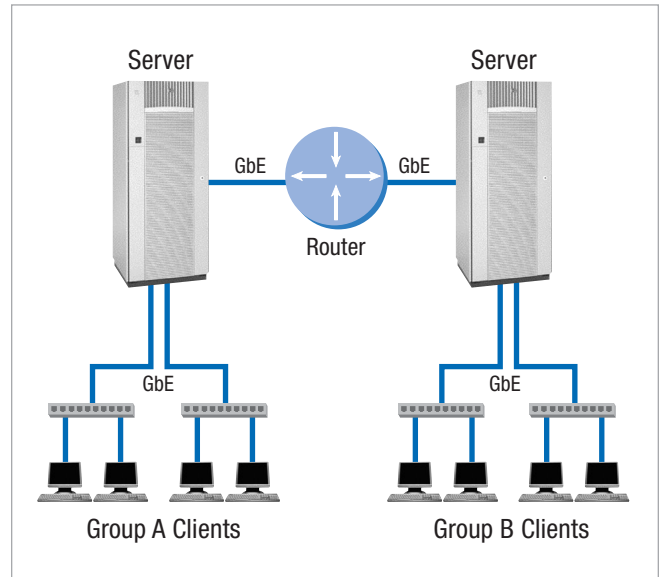


Continuing with the discussion of Figure 1, notice that distribution between the server and the clients is through a simple hub. This hub is a dumb and inexpensive device where each connection sees the traffic of all other clients, which is not a particularly efficient or secure means of distribution. Also, in the case of ten connections or clients, each client sees only one-tenth the bandwidth of the source. Expanding such a network by daisy chaining another hub to the first causes a similar loss of bandwidth for the added hub and its clients.

A better approach to expanding client count is to segment the network as shown in Figure 2. This allows client count expansion while maintaining client port bandwidth, as opposed to diminishing it with daisy chaining. With a static client count, bandwidth and throughput actually increases. This is because now, instead of 10 clients competing for the throughput of a single GbE trunk, segmentation has split the client traffic across two GbE segments, or subnets.

Implementing the segmentation shown in Figure 2 does require an additional GbE server port. Installing another server adapter provides the necessary additional port, but this consumes another server Peripheral Component Interconnect (PCI) or Peripheral Component Interconnect Extended (PCI-X) slot. A better approach when upgrading servers to GbE capability is to use multi-port server adapters. For example, upgrading to GbE capability with a dual-port Intel PRO/1000 Server Adapter uses a single server slot to provide the two GbE ports necessary for the segmentation shown in Figure 2. However,

Figure 3. Segmentation by workgroups

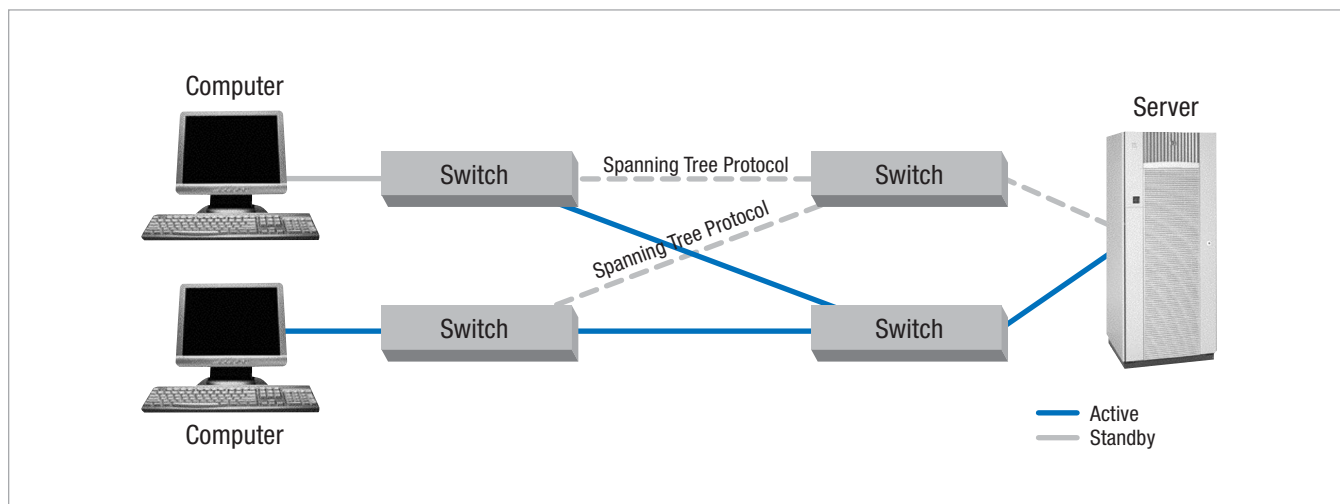


using an Intel® PRO/1000 Quad Port Server Adapter is even better because it provides four GbE ports from a single PCI or PCI-X slot. This conserves server slots while providing ready scalability for additional network growth such as that illustrated in Figure 3.

Figure 3 shows the next evolutionary step in segmenting a small office network, such as might be found in a branch office or a consulting firm. This configuration includes two servers and each server uses three GbE ports. If these servers were initially upgraded with dual-port adapters, they would now have to be upgraded with another adapter to support the implementation of Figure 3. An initial upgrade to a quad-port server adapter, as suggested previously, provides the needed future capacity for transitioning to the segmentation scheme shown in Figure 3.

The segmentation scheme in Figure 3 accommodates the diverse needs of two different workgroups. For example, Group A might represent accounting and management functions, and Group B might represent an engineering function. The segmentation in Figure 3 confines engineering applications and associated traffic to the Group B server. Similarly, management and accounting traffic is confined to the Group A side of the network. Both servers and their associated subnets experience reduced traffic because accounting rarely accesses the engineering side, and engineering rarely accesses the accounting side. They are on different freeways, so to speak. However, both sides still have access to each other via the router for e-mail, budget reports and other cross-enterprise activities.

Figure 4. Switch fault tolerance on a single segment



Such a segmentation approach makes network management easier and increases system reliability and security. For example, because Groups A and B are on different servers, Group B applications or upgrades are less likely to step on Group A applications, and vice versa. Also, if a server bogs down or crashes, it affects only that server's segment of the network. In growing organizations with mission-critical applications, such potential server problems are typically addressed by using server clusters, where the cluster is set up to absorb the loss of any one server.

Switches Enhance Control and Reliability

The segmentation examples shown in Figures 2 and 3 can benefit further by substituting smart switches for the dumb hubs. Dumb hubs distribute network traffic by broadcasting all packets to all client nodes. This forces each node to process all packets in order to find its intended packets. Smart switches minimize this packet processing by routing packets to the intended node only. As a result, packet processing is reduced and throughput increased. The fact that each node receives only the packets intended for it also makes the switched network inherently more secure.

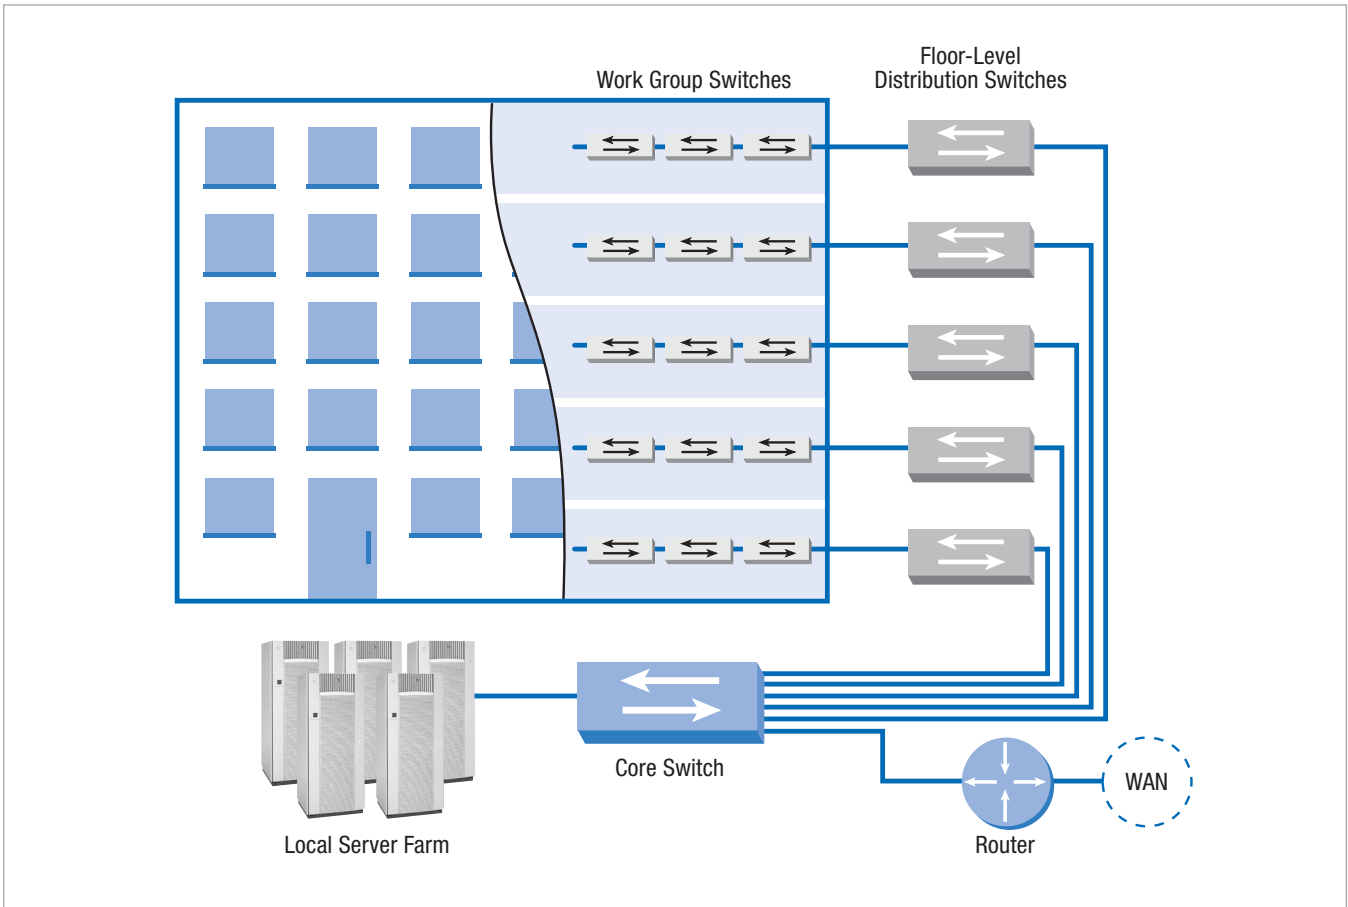
Smart switches can also be configured to make networks more reliable and resilient. Figure 4 illustrates this with multiple switches using spanning tree protocol to provide switch fault tolerance. In this configuration, a sensed failure of a server adapter or link connection results in automatic switching to a secondary standby path (shown by dotted lines in Figure 4).

For simplicity of illustration, Figure 4 shows a single segment on the server. Additional segments with switch fault tolerance require additional sets of cross-connected switches and additional server-adapter pairs. While this does add to network complexity, the benefits are higher throughput through segmentation and higher reliability through switch fault tolerance.

The example of Figure 4 also emphasizes the importance of multi-port server adapters such as the Intel® PRO/1000 MT Quad Port Server Adapter. Using this adapter allows the addition of four GbE ports to a server from a single PCI or PCI-X slot. Slot conservation is particularly important in slot-constrained servers where Redundant Array of Independent Disks (RAID) and other devices are required.

Multi-port server adapters and smart switches become indispensable when implementing larger networks such as the campus Local Area Network (LAN) shown in Figure 5 on the next page. Segmentation in Figure 5 occurs at the local server farm. The servers use multiple GbE ports for segmentation, with the ports applied singly or in pairs for each segment. A hierarchy of smart switches provides distribution to the campus LAN and to each segment within the campus LAN. Depending on organizational and workgroup needs, segments can be arranged on a floor-by-floor basis, or a workgroup segment can span more than one floor.

Figure 5. Example of a campus LAN



Combining NICs and LOM for Performance

As shown in Figures 2, 3 and 4, segmentation requires multiple GbE server ports. Even the simplest case of two segments requires two server ports. Add switch fault tolerance and the port-count requirement doubles. Involving multiple servers requires additional server-to-server or server-to-router ports. Such increased GbE port-count requirements can quickly clash with server slot limits and the demands on server slot inventory for other application cards.

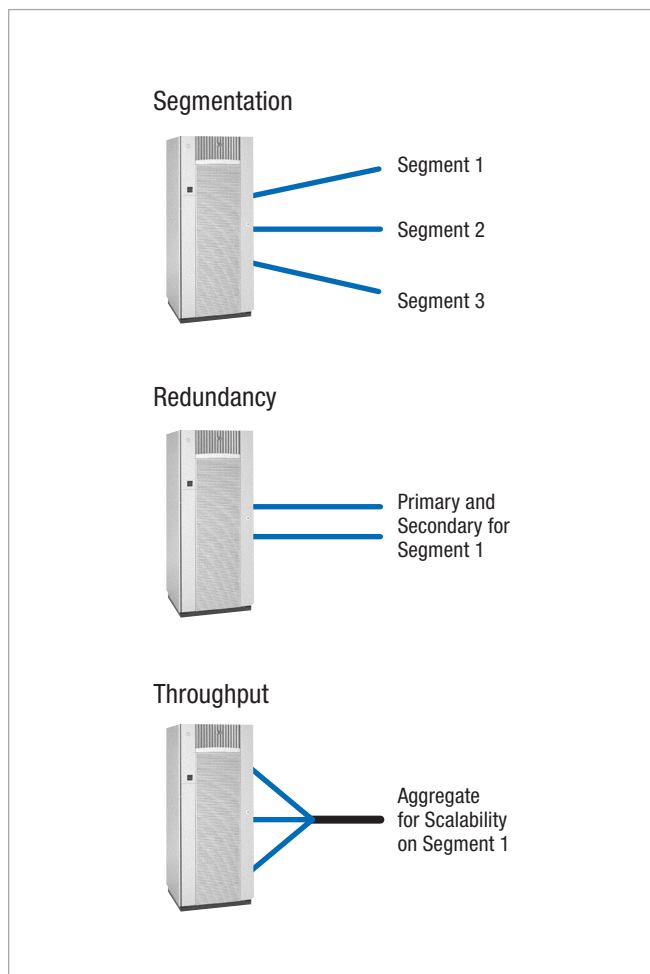
Fortunately, a twofold solution is available for adding GbE server ports while conserving server slots. The first part of the solution is to use servers with embedded LAN on Motherboard (LOM). The second part is to use multi-port Network Interface Cards (NICs), also referred to as server adapters.

LOM provides Ethernet connection without consuming server slots. The ideal situation is that of newer servers offering options for dual-port GbE LOM. This provides both an upgrade to GbE capability while also providing

the minimum two ports necessary for basic network segmentation. However, any further scaling up of network capacity requires additional ports, which must be provided by adding NICs to the server via its PCI or PCI-X slots. You can also use NICs to upgrade older Fast Ethernet servers to GbE capability.

For upgrading to GbE and increasing port counts, Intel offers a full line of single- and multi-port NICs for both wire cable and fiber-optic cable. The Intel PRO/1000 MT Quad Port Server Adapter, for example, uses a single server slot to provide four GbE ports. If the server also has a dual-port GbE LOM, use of the Intel PRO/1000 Quad Port Server Adapter jumps port count to six GbE ports. For servers without LOM, adding two quad-port adapters provides a port count of eight while using only two server slots.

Figure 6. Multiple port scalability examples



As a guideline, there should be no more than three network segments per server. This requires, at minimum, three network ports. However, as shown in Figure 6, additional GbE ports can provide additional increases in network performance. For example, providing a redundant link for each segment increases network reliability by providing an alternate, fail-over path in the event of a connection or link failure. Additionally, higher segment and network throughput can be achieved by using adapter-teaming and link-aggregation techniques.

Redundancy and adapter teaming require additional management software. However, Intel meets that NIC-management need with Intel® Advanced Network Services (ANS), which is a free software download available with Intel PRO/1000 Server Adapters. Intel ANS software

allows you to configure links in a variety of ways for increased performance and reliability. This includes redundancy management and throughput enhancement through:

- **Adapter Fault Tolerance:** Automatically senses failure of the primary adapter and switches traffic to the secondary adapter
- **Link Aggregation:** Increases bandwidth and throughput by aggregating multiple adapters/links into a single virtual link
- **Adaptive Load Balancing:** Automatically senses and balances traffic loads across adapter/link teams

In addition to providing the above management features, Intel ANS software and Intel server adapters are compatible with NICs from other vendors and NICs of differing speeds. This allows you to upgrade network performance and reliability while preserving existing hardware investments.

Further Performance Enhancements

Users are always concerned about performance at the desktop, as well. A high-performance network is difficult to appreciate from a low-performance desktop client. As a consequence, many users are upgrading to GbE network connections in their desktop and notebook connections. This can be done with either new desktop or notebook platforms with GbE LOM options or by upgrading existing client platforms with GbE adapter cards.

With GbE to the desktop, power users gain a performance edge in environments with increasingly complex and file-burdened applications. However, as more desktop clients become GbE capable, GbE network traffic jams will inevitably increase. To counter this, IT managers should consider staying a jump ahead by upgrading server links and network backbones with 10GbE. This can be done with either the Intel® PRO/10GbE LR Server Adapter for single-mode fiber or the Intel® PRO/10GbE SR Server Adapter for short ranges on multi-mode fiber. Upgrading the network core to 10GbE provides the performance margins necessary for riding the upcoming wave of GbE to desktop clients.

Conclusion

Multiple GbE server connections allow IT managers to strategically segment traffic loads to provide higher network capacity, performance and reliability. This can be further augmented by conversion of performance-critical backbones or links with 10GbE capability.

Intel, with its successful history of experience in network connectivity, provides a full range of LOM products and server adapters for multi-port GbE connectivity. In addition to offering GbE capability, these Intel connectivity products contain a variety of features for optimizing throughput and efficiency by unburdening servers from traffic-related processing. These embedded adapter features include Transmission Control Protocol (TCP) Checksum Offload, TCP Segmentation/Large Send Off offload, Interrupt Moderation and auto-negotiation for compatibility with Fast Ethernet clients.

To find out more about Intel® PRO Network Connections, please visit:
www.intel.com/network/connectivity

To find out more about Intel® Advanced Network Services software, please visit:
www.intel.com/network/connectivity/resources/doc_library/white_papers/254031.pdf

This document and related materials and information are provided “as is” with no warranties, express or implied, including but not limited to any implied warranty of merchantability, fitness for a particular purpose, non-infringement of intellectual property rights, or any warranty otherwise arising out of any proposal, specification, or sample. Intel assumes no responsibility for any errors contained in this document and has no liabilities or obligations for any damages arising from or in connection with the use of this document.

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.
Copyright © 2004 Intel Corporation. All rights reserved.

