

Solving Server Bottlenecks and Downtime with Intel® Server Adapters

intel®

Contents

Executive Summary	2
The Problem: Server Bottlenecks	3
Solution: Upgrading to Gigabit Connections	3
Solution: Scalable Server Bandwidth	4
The Problem: Server Downtime	5
Solution: Teaming and Online Serviceability	6
Configuration Considerations using Intel® Server Adapters	6
Reliable, Scalable and Easy to Configure	7

The cost of server downtime is staggering. Less than 1% downtime can cost more than \$3 million per year in lost revenues, user salaries, and server outage costs for a typical medium business[†]. This translates roughly into a loss of \$56,000 per hour of downtime.

[†]Source: Network Server Downtime Cost Analysis, Strategic Research, 1998 (www.networkbuyersguide.com – Online Evaluation Applications). Calculations based on the servers of a 300-employee, \$100M business going down 24 times a year, three hours at a time.

Executive Summary

The proliferation of high-performance desktop and server processors has influenced many LAN managers to migrate their data centers and server links from Ethernet to Fast Ethernet and Gigabit Ethernet. This proliferation continues, driven by the phenomenal growth of the Internet, bandwidth-hungry applications, and powerful workstations. To make the most of their high-speed networks, IT managers are looking for ways to relieve bottlenecks at server connections and other aggregation points.

This technical brief discusses how advanced technologies can address key issues of server performance, including slowdowns due to network connection bottlenecks and downtime due to link failure (see Table 1). These advanced technologies are designed to provide faster connections today, scalability for the future, and automatic failover if a link goes down. All of these can positively impact such critical issues as data center performance and web site response times.

Problems	Causes	Solutions
Slow server performance due to server bottlenecks	<ul style="list-style-type: none"> More desktops running at 100Mbps High-performance processors and faster bus architecture Bandwidth-intensive applications 	<ul style="list-style-type: none"> Upgrading connections to Gigabit Ethernet Scalable server bandwidth, using: <ul style="list-style-type: none"> ■ Adaptive Load Balancing (ALB) ■ Intel® Link Aggregation ■ Fast EtherChannel* (FEC) ■ Gigabit EtherChannel* (GEC) ■ IEEE 802.3ad
Server downtime	<ul style="list-style-type: none"> Failed network connection due to: <ul style="list-style-type: none"> ■ Broken/loose cables ■ Hub or switch port failures ■ Adapter hardware breakdown ■ PCI slot malfunction 	<ul style="list-style-type: none"> Automatic backup links and online serviceability, using: <ul style="list-style-type: none"> ■ Adapter Fault Tolerance (AFT) <ul style="list-style-type: none"> - Mixed-Adapter Teaming - Preferred Primary Technology ■ PCI HotPlug* and Active PCI*

Table 1: Server connectivity problems and their solutions using advanced technologies.

The Problem: Server Bottlenecks

As sophisticated applications and more powerful desktop PCs drive network traffic to new levels, a single 100Mbps channel isn't enough bandwidth for critical server connections – especially as the number of desktops connected at 100Mbps increases.

In the past, such bottlenecks were typically solved by installing an additional Network Interface Card (NIC) in the server, and segmenting the network into two subnetworks (Figure 1). This reduced traffic volume on each network link and eliminated the bottleneck. But segmentation poses a new set of problems, including additional overhead and the need to reassign IP addresses and remap the network. Segmentation generally requires additional hardware, such as switches or routers. Balancing traffic on the two segments can also be difficult, usually requiring repeated reconfiguration. Finally, since the two adapters operate in separate network segments, they don't provide a failover connection in the event of a link failure.

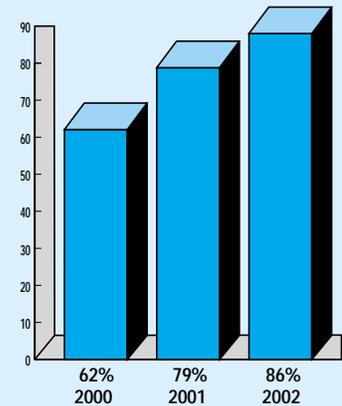
Solution: Upgrading to Gigabit Connections

Since receiving IEEE approval in 1998, Gigabit Ethernet has been used extensively in servers with Gigabit Ethernet network adapters and along backbones to remove traffic bottlenecks in these areas of congestion. In 1999, the IEEE ratified Gigabit Ethernet over copper, allowing speeds up to 1Gbps to be transmitted over Category 5 (Cat-5) cable.

Since Cat-5 makes up a large portion of the installed base, migrating to Gigabit has never been easier. Organizations can now replace 10/100Mbps server adapters with Gigabit adapters, allowing servers to communicate with Gigabit switches at higher speeds, without having to re-wire the infrastructure. Adapters are also available offering Gigabit speeds over fiber cabling for reduced noise over cabling runs of up to 275 meters.

With built-in capability for 10Mbps, 100Mbps and 1000Mbps, adapters such as the Intel® PRO/1000 T Server Adapters allow the network connection to run at 10/100Mbps speeds today while providing a seamless upgrade to Gigabit

10/100Mbps Desktops Running at 100Mbps



Source: Infonetics, 2000

Server Bandwidth – A Growing Need

A 2000 study by Infonetics reveals that 62% of all 10/100Mbps desktops are already running at 100Mbps, and the percentage is rapidly increasing. That same study cited server bottlenecks as the number one reason that businesses are migrating to higher performance networking technologies. Clearly, there's a need for significant increases in server bandwidth to match the growing power and demands of networked clients.

Fortunately, it's possible to significantly increase server bandwidth without a major network overhaul. Upgrading to Gigabit adapters is a good first step, followed by load balancing to increase scalability. Load balancing across multiple Fast Ethernet or Gigabit Ethernet server adapters, using the technologies discussed in this technical brief, provides a simple and scalable solution. And since these load-balancing technologies automatically support redundant network links, they increase server availability as well as performance.

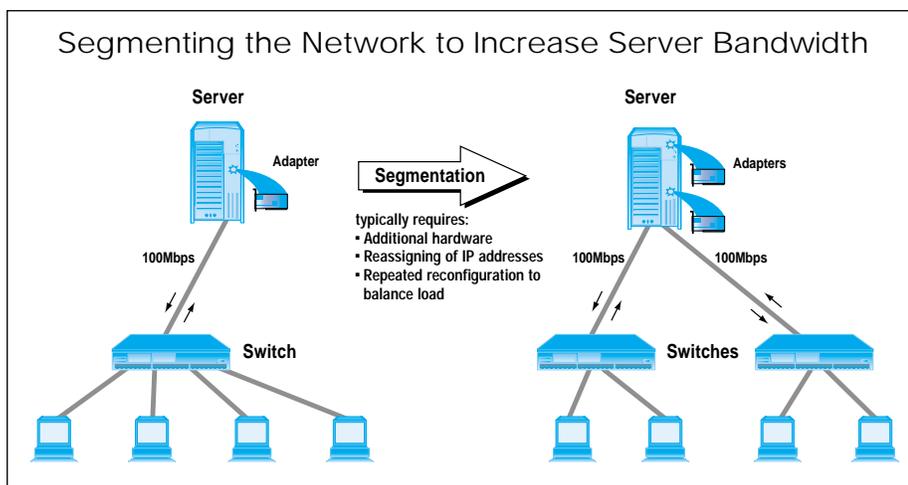


Figure 1: Segmentation increases server bandwidth, but typically requires additional hardware and management overhead – including repeated reconfiguration to balance the traffic load.

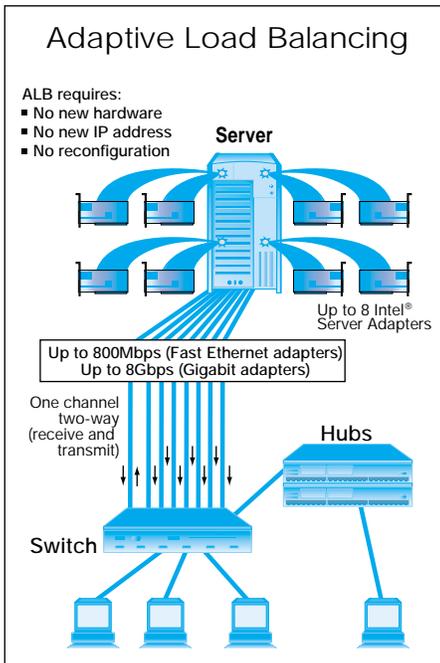


Figure 2: With Adaptive Load Balancing, all traffic traveling from the server is automatically balanced between as many as eight network adapters. This assures fast throughput with no need to restructure or reconfigure the network.

speeds when needed in the future. The Intel® adapter automatically senses when connected switches and routers have been upgraded to Gigabit and will then begin running at Gigabit speeds without needing any intervention.

Solution: Scalable Server Bandwidth

Adaptive Load Balancing

Adaptive Load Balancing (ALB) technology, which can be used with Fast Ethernet or Gigabit Ethernet adapters, offers an additional way to move more data between the server and the network. ALB can increase server bandwidth up to 800Mbps over Fast Ethernet, or 8Gbps over Gigabit Ethernet, by automatically balancing traffic across as many as eight network adapters (Figure 2). Essentially, each additional adapter adds another 100Mbps or 1000Mbps link to the network. Since the distribution of traffic among the adapters is automatic, there is no need to segment or reconfigure the network. The existing IP address of the server is shared by all of the adapters in the server, and traffic is always balanced between them.

ALB is implemented by installing a team of server adapters in the server. The adapters can be quickly configured to run ALB using the Intel® PROSet utility. No client configuration is required, and

clients don't have to be routed to communicate with each other. Moreover, the teamed server adapters provide automatic emergency backup links to the network. If one server link goes down, due to a broken cable, a bad switch port,

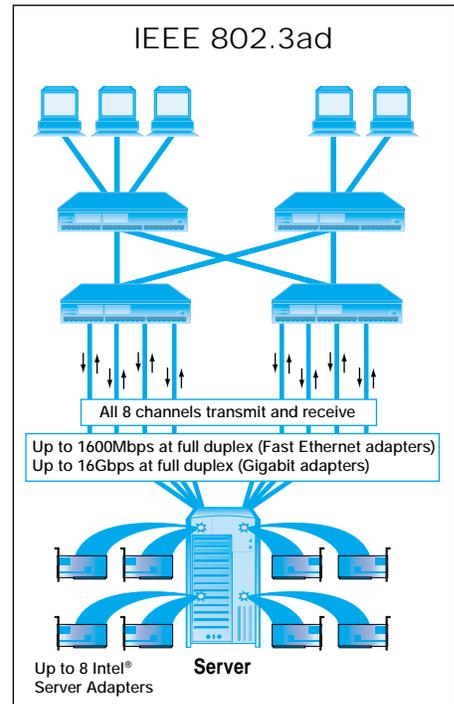


Figure 3: With IEEE 802.3ad link aggregation, network traffic traveling to the server, as well as from the server, is automatically balanced between as many as eight server adapters and multiple switches all on the same network segment.

Advanced Server Adapter Technologies

These technologies provide scalable server bandwidth through load balancing, as well as automatic redundant connections for increased server availability.

Adapter Fault Tolerance (AFT) – Developed by Intel, AFT monitors the server connection to the network and automatically switches traffic to a redundant link in the event of a failure.

Mixed-Adapter Teaming – Enables one kind of server adapter to be used as a redundant backup link for a different kind of server adapter.

Mixed-Speed Teaming – Enables a 100Mbps server adapter to be used as a backup link for a Gigabit server adapter.

PCI HotPlug* and Active PCI* – These standards enable a failed network adapter to be replaced without taking the server offline.

IEEE 802.3ad – This industry standard for link aggregation, ratified in March 2000 and supported by Intel, provides compatibility with other vendors' switches that support the standard. It allows for balancing traffic among multiple server adapters and multiple switches.

Adaptive Load Balancing (ALB) – Developed by Intel, ALB supports scalable bandwidth up to 800Mbps, or 8Gbps in a Gigabit Ethernet environment.

Intel® Link Aggregation – Intel® Link Aggregation supports scalable bandwidth up to 1600Mbps full duplex; or up to 16Gbps in a Gigabit Ethernet environment. Intel Link Aggregation requires support in the NIC and the switch.

Fast EtherChannel* (FEC) – Developed by Cisco, FEC supports scalable bandwidth up to 1600Mbps at full duplex. Fast EtherChannel requires support in the NIC and the switch.

Gigabit EtherChannel* (GEC) – Developed by Cisco, GEC supports scalable bandwidth up to 16Gbps at full duplex. Gigabit EtherChannel requires support in the NIC and the switch.

or a failed adapter, the other adapter(s) automatically accept the additional load (see Adapter Fault Tolerance). There's no interruption in server operation, and a network alert is generated to inform IT staff of the problem.

Up to eight Intel® Server Adapters can be configured to work together as an ALB team. All of the adapters in a team must be connected to a switch. They can be connected to a single switch, or to two or more switches, as long as all the switches are on the same network segment. ALB works with virtually any switch on the market.

Once ALB is configured, all outgoing server traffic will be balanced across the adapter team. Incoming traffic is carried by a single adapter. In most environments this is a highly effective solution, since server traffic is primarily outbound – from the server to the clients.

Intel® Link Aggregation, Fast EtherChannel*, Gigabit EtherChannel*, IEEE 802.3ad

Intel® Link Aggregation and Fast EtherChannel* (FEC) are other technologies that can be used to increase server bandwidth. Like ALB, they automatically balance server traffic and require no network reconfiguration (Figure 3). Unlike ALB, they enable full-duplex transmission on all adapters as long as the switch supports this advanced feature. Both incoming and outgoing server traffic are balanced, and can be scaled in increments. Total throughputs of up to 1600Mbps are possible in Fast Ethernet environments. Intel Link Aggregation can also be used to aggregate traffic across multiple Gigabit server adapters, for throughputs of up to 16Gbps at full duplex. Gigabit EtherChannel*

(GEC) is another technology that provides similar full-duplex load-balancing if connected to supporting GEC switches. A new industry standard for link aggregation, IEEE 802.3ad, allows for balancing traffic among multiple switches from a single server.

Because of their ability to handle high-bandwidth, full-duplex traffic loads, these technologies are ideally suited to high-performance environments running especially demanding applications – such as enterprise servers, web servers, intranet servers and high-end graphics imaging and rendering servers. In addition to scalable server bandwidth, these technologies provide reliable fault tolerance. If one link fails, the other adapters in the team automatically accept the full traffic load, and an alert is generated to notify IT staff of the problem (see Adapter Fault Tolerance).

Whereas ALB works when the Intel Server Adapters are connected to the network via any switch, these full-duplex technologies require that the adapters be connected to switches that support

whichever scalable bandwidth technology is configured in the adapter. Link aggregation is supported by Intel® Gigabit switches and Gigabit uplinks, as well as by an increasing number of other vendors' switch products. Fast EtherChannel and Gigabit EtherChannel work with any FEC- and GEC-enabled switch respectively. Similarly, IEEE 802.3ad requires a capable switch.

The Problem: Server Downtime

Almost every company that depends on networked computers has a nightmare story about server failure. As a result, a variety of mechanisms have been implemented by server manufacturers to improve the reliability of servers. However, a broken or loose network cable, a faulty switch or hub port, or a failed adapter can shut down server operation just as easily as a server malfunction.

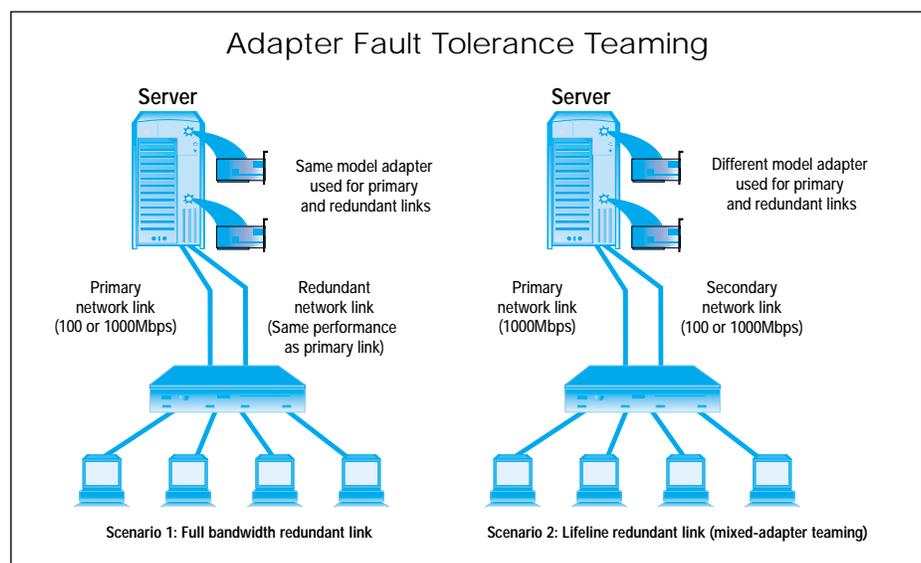


Figure 4: Adapter Fault Tolerance and PCI HotPlug*/Active PCI*.
If the primary network link fails: 1. The failure is immediately detected and the back-up link activated. 2. A network alert is generated to notify IT staff. 3. The problem can be fixed with the server still online. 4. If the failed link was the primary link, it automatically re-establishes itself as the primary link once it is fixed.

Solution: Teaming and Online Serviceability

Adapter Fault Tolerance – Redundant Network Links

Adapter Fault Tolerance (AFT) provides a simple, effective and fail-safe method for avoiding server downtime (Figure 4).

With two or more server adapters installed in a server, AFT can be configured to establish an automatic backup link between the server and the network. Should the primary link fail, the secondary link kicks in within milliseconds, in a manner that is transparent to applications and users.

The redundant link that AFT establishes between the server and the network includes a redundant adapter, a cable, and hub or switch port connection. If there is any problem along the primary link, the secondary link immediately takes over. AFT also initiates a network alert. The server remains online so technicians can take corrective measures when appropriate – during off-business hours, for example.

AFT can be implemented in a server using only two server adapters, one as the primary connection and the second as a backup. AFT is also supported when server adapter teams are configured for ALB, Intel Link Aggregation, Fast EtherChannel, Gigabit EtherChannel or

IEEE 802.3ad. In those cases, if any of the server links fail, for whatever reason, the remaining links automatically take over to share the traffic load.

Mixed-Speed and Preferred Primary Teaming

Unlike most redundant link technologies, AFT supports mixed-speed teaming using any combination of Intel Server Adapters. For example, a Gigabit server adapter could be used as the primary network link. The backup link could be another Gigabit or Fast Ethernet server adapter. With this capability, a relatively inexpensive, 100Mbps backup link can be used to safeguard a high-speed Gigabit Ethernet connection. The inexpensive backup may not be able to support the full traffic load as effectively, but it can allow business critical applications to stay online until the higher speed link is fixed.

When configuring AFT, a preferred primary adapter can be specified. If the primary link fails, it will automatically be reinstated as the primary link once it is fixed. For example, if a Gigabit server adapter is being used for specialized, high-demand applications, a less expensive backup link can be installed using a Fast Ethernet server adapter. If the primary link fails and is then fixed, traffic will automatically revert back to the higher performance link.

PCI HotPlug* and Active PCI*

These technologies enable a failed adapter to be replaced without taking the server offline. Developed by Compaq and IBM, respectively, they have since been established as industry standards, supported in most new servers.

When used in conjunction with AFT, they allow an adapter to be replaced without interrupting network service. If an adapter fails, AFT automatically moves server traffic onto the redundant link and generates a network alert. PCI HotPlug* and Active PCI* enable IT staff to replace the failed adapter without bringing down the server.

Configuration Considerations using Intel® Server Adapters

A single driver provides the software agent that supports AFT, ALB, Intel Link Aggregation, Fast EtherChannel, Gigabit EtherChannel, PCI HotPlug, Active PCI, and IEEE 802.3ad. How the agent is configured in a particular environment determines which of the advanced features is enabled. However, all the scalable bandwidth technologies supported by Intel Server Adapters include built-in support for AFT, PCI HotPlug and Active PCI. So, if ALB, Intel Link Aggregation or Fast EtherChannel is configured, the others are automatically activated.

Intel® Server Adapters – Support for Advanced Features

Intel® Server Adapter	Application	AFT	HotPlug*	ALB	Link Agg.	FEC*	GEC*	802.3ad
Intel® PRO/1000 T Server Adapter	High-traffic backbone servers for Cat-5 copper	✓	✓	✓	✓	✓	✓	✓
Intel® PRO/1000 F Server Adapter	High-traffic backbone servers for fiber	✓	✓	✓	✓	N/A	✓	✓
Intel® PRO/100+ Dual Port Server Adapter	Departmental servers with limited PCI slots	✓	✓	✓	✓	✓	N/A	✓
Intel® PRO/100 S Server Adapter	Departmental or workgroup servers for security	✓	✓	✓	✓	✓	N/A	✓

Table 2: Intel® Server Adapters – support for advanced features: AFT – Adapter Fault Tolerance; HotPlug – PCI HotPlug*; ALB – Adaptive Load Balancing; Link Agg. – Intel® Link Aggregation; FEC – Fast EtherChannel*; GEC – Gigabit EtherChannel*; 802.3ad - IEEE 802.3ad standard.

All the advanced server adapter features supported by Intel Server Adapters integrate seamlessly into Microsoft Windows*, Windows NT*, Linux*, and Novell NetWare* operating system-based servers. The advanced features are management-ready and simple to use, with intuitive interfaces for quick setup and ease-of-use. Standard operating system interfaces are used for Linux and NetWare, while Windows 2000 and Windows NT use PROSet, Intel's intuitive Windows operating system-based configuration utility. Network alerts for failed links are operating system-based for compatibility with management applications. Specifically, NetWare alerts are generated for NetWare servers and event logs for Windows NT and Windows 2000 servers. A management application can detect these alerts and trigger an appropriate action. For example, a network manager could choose to be notified of a failure via an e-mail

message, a fax or a call to his pager or cellular phone.

Intel's newest generation of PRO Server Adapters are backwards compatible with previous generation Intel® PRO Server Adapters based on Intel® 82558, 82559, 82550 technology, as well as Intel® 82542 and 82543 LOM and network adapters based on this silicon. This means that an organization's investment in these Intel® products can be preserved by using them in a server adapter team to increase redundancy and scale bandwidth.

Reliable, Scalable and Easy to Configure

Scalable bandwidth technologies, along with AFT, PCI HotPlug and Active PCI make the Intel Server Adapter family an ideal solution for fast network connectivity with enhanced server availability. Table 2 shows which advanced features are supported by which

Intel Server Adapters. In practice, each business can choose to configure the adapter software to use the technologies best suited to the demands of their server environment and their existing infrastructure. Each technology builds on the preceding one, so nothing is lost as higher bandwidth load-balancing technologies are employed (Table 3).

By providing scalable bandwidth and increased availability at a crucial point in the network, Intel Server Adapters can help to revive network infrastructures that are otherwise straining under increased traffic loads. They also enable a more affordable server solution for high-demand networks. By integrating high-availability server links and load balancing into the network adapter, they eliminate the need for specialized server hardware and other expensive infrastructure components.

Advanced Feature	Connection Requirements	Benefits
Adapter Fault Tolerance with: – Mixed-Speed Teaming – Preferred Primary	When connected to any hub or switch	<ul style="list-style-type: none"> High Server Availability – Transparent backup connections using any combination of Intel® Server Adapters
PCI HotPlug*/Active PCI*	When connected to any hub or switch and supported in the server	<ul style="list-style-type: none"> Online Serviceability – Adapters can be replaced with the server online
Adaptive Load Balancing	When connected to any switch	<ul style="list-style-type: none"> High Server Availability Load balancing of outgoing server traffic for increased throughputs up to 800Mbps (8Gbps 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	<ul style="list-style-type: none"> High Server Availability Load balancing of outgoing and incoming server traffic, for throughputs up to 1600Mbps (16Gbps in Gigabit Ethernet environments)
IEEE 802.3ad	When connected to a switch that supports IEEE 802.3ad link aggregation	<ul style="list-style-type: none"> Higher Server Availability Balances traffic among multiple switches from one server for higher throughput and increased redundancy

Table 3: By configuring the adapter for the desired advanced feature support, companies can match the needs and resources of their particular server environment.

As the number one supplier of Fast Ethernet and Gigabit Ethernet connectivity solutions[†], Intel has played a leading role in the development of advanced technologies for server connectivity.

For more information on Intel® Server Adapters:

- Product information: http://www.intel.com/network/products/server_adapters.htm
- NOS support for advanced server adapter technologies: http://www.intel.com/network/technologies/advanced_features.htm

Related Intel® White Papers:

- Layer 2 Network Prioritization: http://www.intel.com/network/white_papers/priority_packet.htm
- Building a Managed Computing Environment: http://www.intel.com/network/white_papers/managed_environment.htm
- Gigabit Ethernet Solutions: http://www.intel.com/network/solutions/ethernet_guide/index.htm
- Gigabit Ethernet over Copper: http://www.intel.com/network/solutions/copper_guide/index.htm

[†] Source: Dell'Oro Group, November 2000

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice. For the most current product information, please visit <http://www.intel.com/network>