

# Gigabit Networking Gigabit Ethernet Solutions

As the leading provider of switched internetworking solutions, Cisco Systems is committed to the development of technology and products that provide gigabit transmission speeds for enterprise networks. Cisco is investing aggressively in Gigabit Ethernet technology and products and will deliver complete networking solutions to scale campus intranets.

This white paper specifically addresses the following:

- The Situation Creating the Evolution  
Toward Gigabit Networking
- Key Challenges Facing the Industry Evolution  
To Gigabit Networking
- Gigabit Ethernet Establishes a Foundation for  
Gigabit Networking
- Cisco Meets Industry Challenges with Complete Solutions
- Cisco Scales Technology Elements Required  
for Gigabit Networking
- Campus Networks Designed to Scale
- Cisco's Path to Gigabit Multilayer Switching  
for the Campus

## The Situation Creating the Evolution Toward Gigabit Networking

Several industry trends are leading enterprise users to examine the need for gigabit networks. Each company situation will be different, requiring a specific set of migration steps to handle growing network traffic and changing traffic patterns. Key industry trends creating the evolution toward gigabit networking include the emergence of the intranet network model, higher bandwidths required by network users, increases in processor power, new applications, and changing traffic patterns.

## The Intranet Network Model

Without question, today's most significant drivers within the enterprise network are Internet, intranet, and extranet technologies. This is reflected in the explosive use of Web technologies that are fundamentally transforming the manner in which business is conducted. In addition, a higher number of network users depend on traditional applications (such as file transfer, e-mail, and network backup) to conduct business, creating a steady growth of network traffic. The result is a geometric growth in traffic and a permanent change in the nature of enterprise networks, as well as an increase in the commercial assimilation of both protocols and application styles that began life on the Internet.

One of the most significant changes that has occurred is the unpredictable network traffic patterns that result from the combination of intranet traffic, fewer centralized campus server locations, and the increasing use of multicast applications. The old 80/20 rule, which stated that only 20 percent of network traffic went over the backbone, has been scrapped. The ease with which internal Web browsing now enables users to locate and access information anywhere on the corporate intranet means that traffic patterns are dictated by where the servers with the most valuable pages are and not by the physical workgroup configurations with which they happen to be grouped. Thus the vast majority of traffic will traverse the backbone, and any-to-any traffic will become the rule. Simplified access has also sent overall usage of the network skyrocketing, as users point and click their way through the corporate portfolio of Web-based resources.

Without question, intranets are a key design center for the implementation of sophisticated multimedia application styles that are becoming increasingly complemented by the use of infochannel push technology. Complement this trend by the doubling of uniprocessor power and the quadrupling of multiprocessor power every eighteen months, and you have the basis for nearly geometric growth in network performance requirements that are being generated on a year-by-year basis.

The long-term implication is that users will increasingly demand solutions for *gigabit networking* whose specific requirements include tens if not hundreds of gigabits per second of total network capacity, any-to-any communication, smooth migration of scalable performance that can be incrementally implemented anywhere in the network, and strong compatibility with the infrastructure of existing enterprise networks.

#### **Network Users Require Higher Performance with Smooth Migration**

Growing numbers of users on the network, more current applications, faster desktop computers, and faster network servers create a demand for higher-performance LAN segment capacity and faster response times. Bandwidth enhancement beyond Fast Ethernet is needed to provide smooth network operation in the face of emerging bandwidth requirements. Improvements in network-layer performance to hundreds of thousands of packets per second or millions of packets per second is required to meet the challenge of high-performance network-layer throughput and changing traffic patterns.

New users, new applications (such as multimedia, Internet access, and groupware), and new high-performance servers that are centralized all contribute to traffic congestion and changing traffic patterns on the backbone. Migration toward gigabit bandwidth on the LAN backbone will provide the infrastructure to meet the needs of evolving enterprise networks.

Preservation of installed user applications, network operating systems, network equipment, and network management are highly desirable, as network bandwidth migrates to gigabits per second. However, the increase in bandwidth and network-layer performance must be available in incremental, manageable steps; thus providing a migration to gigabit networking, which is cost effective and practical to implement.

#### **Computing Products Increase Processor, Network, and Input/Output Performance**

Desktop CPUs are rapidly increasing in speed. The Peripheral Component Interconnect (PCI) bus is becoming increasingly popular on many computing platforms. This high-performance bus enables desktop CPUs to fully utilize the bandwidth of Fast Ethernet connections. Manufacturers of high-end desktop systems are committed to offer 100-Mbps Ethernet connections on their motherboards. Workstation and server technology is advancing to enable CPUs to flood multiple Fast Ethernet links with network traffic. Each of these technology trends points to the need for gigabit networking technology that can be deployed for backbone, server, and, eventually, desktop connections.

Uniprocessor performance doubles every 18 months. Multiprocessing system performance increases by a factor of 4 every 18 months. Network technology improvements are needed to keep up with these advances.

#### **New Applications and More Users Demand More Bandwidth from the Network**

PCs, servers, and networks have become key elements of today's business operations. The growing number of users increases traffic on the network. Most enterprises report 200 to 300 percent traffic growth per year as a result of an increasing number of users and higher dependence on existing applications.

New applications such as Web-based intranets, whiteboarding, and desktop video are growing rapidly. Enterprises with intranet-centric business operations today see traffic growth at 500 to 600 percent per year. This is in addition to the increasing number of users. Thus, there needs to be a solution for scaling the campus intranet to handle this new traffic.



## Campus Traffic Patterns Move Toward Any-to-Any Communication

Factors contributing to traffic pattern changes in campus networks include the application of Web-based technologies, centralization of servers in a few locations on the campus, and the adoption of multicast applications. The adoption of Web-based technologies for campus intranets is proving to be a powerful productivity tool. However, the Web-based tools enable access to files located anywhere on the campus and even on the Internet. With the simple click of a mouse, files are summoned from multiple campus server locations to build a Web browser page at the desktop. As Web-based technology is adopted for productivity improvement and information sharing within corporations, traffic patterns shift toward anywhere-to-anywhere, stressing network-layer performance.

The centralization of campus servers provides network managers with benefits in ease of configuration, control, and management. Centralization also enables justification of high-performance server systems for larger numbers of users. As servers are centralized, all traffic must cross the backbone to reach the servers and return to desktops. Such traffic patterns stress the building and campus backbone bandwidth, as well as the network layer. Migration to improved performance will be required.

Multicast applications generate network traffic based on who needs the information and joins the multicast group. Membership can come from anywhere within the campus or enterprise network. Applications such as desktop video, file sharing, and whiteboarding will influence traffic patterns through constant changes in team members. As team members enter and leave virtual work teams, changing the multicast group membership, the multicast traffic patterns will change creating dynamic anywhere-to-anywhere traffic.

## Key Challenges Facing the Industry Evolution To Gigabit Networking

The emergence of Gigabit Ethernet creates a Media Access Control (MAC) and physical sublayer (PHY) standard with the potential for wide deployment of low-cost gigabit campus backbones, server connections, and, eventually, desktop connections. Thus, Gigabit Ethernet will become a standard interface for gigabit networking equipment. However, the deployment of gigabit networking creates a number of challenges beyond the MAC and PHY interfaces. These challenges include multigigabit system bandwidth, Layer 3 forwarding and routing at gigabit line rates, application of network services at gigabit line rates, monitoring and management for gigabit systems, and a smooth migration path from today's networks to gigabit networking.

The deployment of Fast Ethernet products provides a guide for the expected deployment of Gigabit Ethernet. After the Fast Ethernet MAC and PHY standard was finalized, the first products to market were 100-Mbps uplinks for Ethernet switches, 100-Mbps server adapters, and router interfaces. Gigabit Ethernet products will follow a similar pattern. After the IEEE 802.3z standard for Gigabit Ethernet is completed, the implementation will be straightforward for Gigabit Ethernet MAC and PHY functions. These functions will be connected into existing Fast Ethernet switching fabrics, implemented on Gigabit Ethernet router interfaces, and used to create Gigabit Ethernet server adapters. Thus, the first Gigabit Ethernet products to market will be uplinks for Fast Ethernet switches, Gigabit Ethernet server adapters, and uplinks for existing

routers. Beyond these initial Gigabit Ethernet products, there are several challenges that must be addressed to enable the full functionality of gigabit networking.

Key challenges facing the industry include:

- Multigigabit system bandwidth—An architecture and robust system implementation must be delivered that can switch multigigabit data streams and scale to meet the requirements of growing campus intranets.
- Network-layer forwarding and routing at gigabit line rates—Network-layer throughput at gigabit rates, incorporation of multiprotocol capability, operation of routing protocols that support gigabit rate forwarding, and delivery of network services are required to fully utilize the bandwidth of Gigabit Ethernet.
- Application of network services at gigabit rates—Network services such as security, quality of service (QoS), resiliency, statistics, and policy implementation must be supported at gigabit rates. Each packet or flow must be analyzed for application of network-layer services. This procedure requires analysis, decisions, and application of services at very high speed, with a reliable and proven implementation.
- Monitoring and management of gigabit systems—Gigabit rates demand new approaches to monitoring, troubleshooting, and managing of multigigabit systems. While Remote Monitoring (RMON) I and II provide the basis for going forward, unique application of these tools is necessary to process frames that are fast enough for gigabit systems at the core of campus intranets.
- Smooth, scalable migration—Gigabit Ethernet switching and interface products must deliver a smooth, stable, and scalable migration. There must be simple, straightforward steps for migrating performance, capacity, and network-layer functionality, while the system maintains network resiliency for mission-critical applications. Coexistence and interoperability with ATM campus, metropolitan, and wide-area technologies is also required. The introduction

of Gigabit Ethernet switching requires integration into mission-critical networks that have already deployed a backbone technology while maintaining network integrity.

## **Gigabit Ethernet Establishes a Foundation for Gigabit Networking**

### **Gigabit Ethernet Is a Natural Upgrade Path**

The growing importance of LANs today and the increasing complexity of desktop computing applications are fueling the need for high-speed networks. The bandwidth provided by a 10-Mbps Ethernet connection may not be an adequate match for today's typical desktop computing applications.

Numerous high-speed LAN technologies have been proposed to provide greater bandwidth and improved client/server response times. Foremost among them is Fast Ethernet, or 100BaseT, a technology designed to provide a nondisruptive, smooth evolution from 10BaseT Ethernet to high-speed 100-Mbps performance. Given the trend toward 100BaseT connections to the desktop, the need for even higher-speed connections at the server and backbone level is clear.

Gigabit Ethernet will be ideal for deployment as a backbone interconnect between 10/100BaseT switches and as a connection to high-performance servers. A natural upgrade path for future high-end desktop computers, Gigabit Ethernet will require more bandwidth than can be provided by 100BaseT.

### **The Emergence of Gigabit Ethernet**

The IEEE 802.3z Task Force has issued a draft for Gigabit Ethernet. The draft allows for both half- and full-duplex operation with a variety of physical interfaces. Thus, there will be switched and shared topology implementations for Gigabit Ethernet (see Table 1). The choice of topology will depend on the network connection objectives. For example, switched topologies provide the longest distance and high throughput. Shared topologies will provide lower cost with shorter distance capabilities. Cisco is developing


products that are targeted for switched topologies and that will be compatible with interfaces that operate in shared environment (half-duplex mode).

Table 1 Gigabit Ethernet Topologies

Topology	Objective	Modes	Media	Connection Applications
Switched	<ul style="list-style-type: none"> <li>High throughput</li> <li>Long distance</li> </ul>	<ul style="list-style-type: none"> <li>Full duplex</li> <li>Half duplex</li> </ul>	<ul style="list-style-type: none"> <li>Multimode</li> <li>Single-mode</li> <li>Copper</li> </ul>	<ul style="list-style-type: none"> <li>Campus backbone</li> <li>Building backbone</li> <li>Wiring closet uplinks</li> <li>Servers</li> </ul>
Shared	<ul style="list-style-type: none"> <li>Low cost</li> <li>Short distance</li> </ul>	<ul style="list-style-type: none"> <li>Half duplex "Classic repeater"</li> </ul>	<ul style="list-style-type: none"> <li>Multimode</li> <li>Copper</li> </ul>	<ul style="list-style-type: none"> <li>Servers</li> <li>Desktops (long term)</li> </ul>
	<ul style="list-style-type: none"> <li>Low cost</li> <li>Long distance</li> </ul>	<ul style="list-style-type: none"> <li>Full duplex "Buffered distributor"</li> </ul>	<ul style="list-style-type: none"> <li>Multimode</li> <li>Copper</li> </ul>	<ul style="list-style-type: none"> <li>Servers</li> <li>Desktops (long term)</li> </ul>

A primary goal of the IEEE 802.3z Task Force, which is the body developing the standard for Gigabit Ethernet, is to be compatible with installed LAN media: single-mode fiber, multimode fiber, and balanced shielded copper cable. The draft of the IEEE 802.3z specifies three physical interfaces: 1000BaseLX, 1000BaseSX, 1000BaseCX for initial Gigabit Ethernet implementations. The IEEE 802.3ab Task Force will specify a fourth physical interface for Gigabit Ethernet over Category 5 unshielded twisted-pair (UTP) on a separate timetable (9 to 12 months behind the 802.3z schedule [see Figure 2]). Figure 1 summarizes the interfaces

Figure 1 Gigabit Ethernet Media Distances



and link distances for Gigabit Ethernet. All distances in the chart assume a switched interface operating in full-duplex mode and are based on the IEEE 802.3z draft. Due to the collision timing constraints, half-duplex media using carrier sense multiple access collision detect (CSMA/CD) is limited to 100-meters link length, regardless of media).

#### Gigabit Ethernet Standard Time Line

The IEEE 802.3z Task Force is responsible for delivering a finalized standard that meets the objectives established by the IEEE 802.3 executive committee. Cisco is leading the standards activity through the chairmanship of the IEEE 802.3z Task Force and by actively contributing to the technical content of the draft standard documents.

The IEEE 802.3z Task Force has established an aggressive timetable for development of the Gigabit Ethernet standard. In July 1996, the 802.3z Task Force was established with the charter to develop a standard for Gigabit Ethernet. The goals of the IEEE 802.3z Task Force include: use of the 802.3 frame format, half- and full-duplex operation, use of CSMA/CD, optional flow control with 802.3x, and backward compatibility with installed media (single-mode fiber, multimode fiber, UTP copper). Basic concept agreement on technical contributions for the standard was achieved at the November 1996 IEEE meeting. A draft of the IEEE 802.3z standard is now under review and comment by the IEEE 802.3z Task Force. In July 1997, the 802.3z working group ballot began. This is the final internal checkpoint for technical content and comment resolution before the standard goes out for public review and comment in late 1997. After public review and comment resolution, the final ballot is expected in Q1 CY'98.

The final version of IEEE 802.3z is expected to include any changes found to be necessary based on interoperability test results conducted by the Gigabit Ethernet Interoperability Test Consortium in late 1997 or early 1998. The Test Consortium has established an interoperability test lab at the University of New Hampshire to conduct vendor- neutral interoperability testing based on specific test suites designed to test conformance with the IEEE 802.3z specification.

While this schedule is aggressive, it is very similar to the Fast Ethernet standard. The Gigabit Ethernet standard is being developed by the same individuals and companies who developed Ethernet and Fast Ethernet. These individuals are very familiar with Ethernet technology, the standards process, and the requirements of the IEEE to achieve standardization.

Product implementation will be based on the IEEE 802.3z standard time line. Depending on the time of shipment, products from any vendor will fall into one of three Gigabit Ethernet categories: prestandard, interoperable, and production.


**Prestandard**—Any Gigabit Ethernet product design that was finalized prior to the first draft IEEE 802.3z (January 1997) is a “guess” and potentially at risk for interoperability. The IEEE 802.3z draft has gone through three revisions since January 1997, and many vendors have not kept up with the specification changes. Any products shipping in the first half of 1997 will fall into this prestandard category and should serve as a “red flag” to users (see Figure 2). In fact, any products shipping before the finalization of IEEE 802.3z are prestandard.

**Interoperable**—Now that a draft of the standard exists (January to May 1997), network equipment suppliers will be able to develop interoperable products compliant with the draft in the second half of 1997. The Gigabit Ethernet Test Consortium is developing test suites for interoperability testing and is expected to have these complete by mid-October 1997. The test suites will form the basis for the first vendor-neutral interoperability testing of Gigabit Ethernet interfaces in late November or early December. Products shipping during H2 CY’97 that are not compliant to the final standard will require field upgrade. This should serve as a “yellow flag” to users (see Figure 2).

**Production**—The working group ballot milestone for the IEEE 802.3z Task Force is expected to be completed in Fall 1997. Completion of this milestone indicates that the internal 802.3z review of the standard is complete, and only incorporation of interoperability test results and public review remain. The working group ballot milestone indicates a high degree of confidence in the technical aspects of the draft standard. Network equipment vendors are now able to implement product designs with high confidence that products will be fully standards compliant and interoperable. This leads to production-worthy Gigabit Ethernet products in the first half of 1998.

The bottom line—1998 is the year for initial Gigabit Ethernet production product deployment. Cisco is investing aggressively in Gigabit Ethernet technology and product development and will be compliant with the final IEEE 802.3z standard. This compliance assures Cisco customers that they will not have to replace prestandard products after the IEEE 802.3z standard is finalized.

Figure 2 Gigabit Ethernet Time Line




### Gigabit Ethernet for Campus Intranet Applications

One application for Gigabit Ethernet is the building backbone. For this application, Gigabit Ethernet is deployed for backbone links in the building riser that connects a centrally located switch with each wiring closet. Each wiring closet switch has a Gigabit Ethernet uplink. Multimode or single-mode media is used to achieve the required distance. A Gigabit Ethernet switch is centrally located in the building data center with connection to servers, routers, and Asynchronous Transfer Mode (ATM) switches as needed. The server connections can use copper or short-distance fiber for lower cost. Routing and ATM services are provided as needed for high-speed connection to the wide-area network (WAN) or metropolitan-area network (MAN).

A second application for Gigabit Ethernet is the campus backbone. Here, Gigabit Ethernet links are used to connect switches in each building with a central campus switch. Full-duplex operation achieves maximum throughput and distance with fiber media. Either single-mode or multimode fiber can be used. A Gigabit Ethernet switch is located in a central location with connection to servers, routers, and ATM switches as needed. The server connections within the campus data center can use copper or short-distance fiber for lower cost. Routing and ATM services are provided as needed from the campus data center for high-speed connection to the WAN (see Figure 3).

Figure 3 Building and Campus Applications of Gigabit Ethernet



Gigabit Ethernet is well suited for connecting high-performance servers to the network. Servers are growing in power and throughput. Processing power is doubling every 18 months. This growth, combined with the trend for centralizing servers within large enterprises, results in a requirement for very-high-bandwidth network connections. Today, high-performance UNIX servers are able to flood three to four Fast Ethernet connections simultaneously. As the processing power of these systems grows, they will require a faster network. Gigabit Ethernet is ideally suited to provide the high-speed network connection.

In specific applications such as animation, film postproduction, or image processing that require transfer of larger files between desktops and servers, network performance is directly proportional to business productivity. Gigabit Ethernet will provide a solution to current network performance constraints. In the short term, Gigabit Ethernet is expected to be deployed for backbone and server connections. As desktop systems continue to grow in power and as applications require more bandwidth and faster response time from the network, Gigabit Ethernet will be deployed at the desktop to meet these specific requirements.

#### **Gigabit Ethernet Alliance Promotes Industry Cooperation**

The Gigabit Ethernet Alliance is an open forum dedicated to promoting industry cooperation in the development of Gigabit Ethernet. Cisco is a leading contributor to the Gigabit Ethernet Alliance objectives and initiatives. Cisco is leading the Alliance activities through the vice-chairmanship and regularly contributes to initiatives that support Alliance objectives.

The Alliance's primary objectives include:

- Full support of the Gigabit Ethernet standards activities being conducted in the IEEE 802.3 working group
- Contribution of technical resources to facilitate convergence and consensus on technical specifications
- Provision of resources to establish and demonstrate product interoperability
- Fostering of two-way communication between potential suppliers and consumers of Gigabit Ethernet products

The Gigabit Ethernet Alliance builds on its members' past experience and success with the Fast Ethernet Alliance. The Gigabit Ethernet Alliance was founded by Bay Networks, Cisco Systems, Compaq, Granite Systems, Intel, LSI Logic, Packet Engines, 3Com, Sun Microsystems, UB Networks, and VLSI. It has the following organizational structure:

- A steering committee, which is responsible for oversight of all Alliance activities
- A technical subgroup
- A marketing and communications subgroup

Today, the membership has grown to more than 110 companies that represent the networking, computer, semiconductor, test equipment, and component industries. Membership in the Alliance and participation in Alliance activities are open to all interested parties. For more information, visit the [www.gigabit-ethernet.org](http://www.gigabit-ethernet.org) Web page.

#### **Gigabit Ethernet and ATM are Complementary for Campus Intranets**

Criteria for Scaling Campus Intranets

The emergence of Gigabit Ethernet creates the option to choose between two high-speed LAN technologies—ATM and Gigabit Ethernet—for campus backbone, building backbone, and server applications. The choice of technology for scaling campus intranets is best accomplished by establishing a set of criteria that can be viewed as a set of operational goals that deliver a scalable infrastructure that meets the mission-critical requirements of the business. For example, the following criteria can be used to evaluate Gigabit Ethernet and ATM technology for use in scaling the performance of your campus intranet.

To successfully scale a campus network or intranet, several key requirements must be considered:

- Bandwidth performance/latency performance—Existing and emerging applications will require higher bandwidth and lower latency. Traffic patterns are changing, are less predictable, and require higher bandwidth to handle backbone traffic.
- Compatibility with installed server, desktop, and network equipment—Large enterprise networks have invested millions of dollars for server, desktop, and network infrastructure. Compatibility and leverage of this previous investment is important. Easy migration with the existing infrastructure is also important.

- Compatibility with installed LAN protocols—Protocol compatibility is important to leverage existing applications and to smooth migration.
- Emerging needs for QoS—These emerging needs on the campus must be addressed. Options include guaranteed QoS, less stringent but “good enough” CoS, or the simplest approach: lots of bandwidth.
- WAN compatibility—For enterprises designing their campus backbone, WAN compatibility is significant for scaling the campus network, where WAN traffic and access costs dominate and must be carefully evaluated.
- Services integration—Integration of data, video, and voice services can be a key objective for cost reduction. Here, consolidation of WAN services, the campus backbone, and management simplify and lower cost.
- Product availability—The availability of products is a key factor and must be consistent with the planned deployment for scaling your campus network performance. Products include networking products, embedded management agents, management applications, and test equipment.

Tables 2 and 3 summarize the capabilities for both Gigabit Ethernet and ATM. Note that each technology has its own strengths, and the choice will ultimately depend on the existing network and the role established for the campus network in the future.

There are areas in Table 2 where Ethernet capabilities are increasing. Ethernet is now scalable with 10/100/1000 Mbps. Cisco’s Fast EtherChannel® technology further increases the scalability of Ethernet. Prioritized QoS capability will be available with the finalization of IEEE 802.1Q/p and IETF Resource Reservation Protocol (RSVP) standards. Finally, Packet-over-SONET (POS) enables high-speed, packet-oriented WAN connections without the need to assemble/disassemble packets into cells for transmission over the WAN. The capabilities of Ethernet are increasing toward the current capabilities of ATM while preserving the compatibility with the installed LAN nodes (80 percent of which are Ethernet) and installed protocols (which all operate over Ethernet).

Table 3 Gigabit Ethernet and ATM Feature Comparison

Key Features	Gigabit Ethernet	ATM
Bandwidth	Low Cost	Moderate Cost
QoS	Class of Service (CoS) with 802.1Q, RSVP	Guaranteed QoS
Service Integration	High-Speed Data, Potential, for Voice/Video over IP	Data, Video, Voice
Applications	Building Backbone/Riser Campus Backbone Servers	WAN Building Backbone/Riser Campus Backbone Servers
Product Availability	1998	Shipping Now

Table 2 Ethernet and ATM Comparison with Campus Scaling Criteria

Compatibility	Installed End Stations	LAN Protocols	Scalability	QoS	WAN
Ethernet Packets	Yes	Yes	Yes	Emerging	Emerging
ATM Cells	Yes, with LANE and MPOA	Yes, with LANE and MPOA	Yes, with LANE and MPOA	Yes	Yes

Table 3 summarizes the key features of ATM and Gigabit Ethernet. Notice the differences in bandwidth: Gigabit Ethernet provides low-cost bandwidth, while ATM provides similar bandwidth at higher cost. This is due to the higher level of services and functionality provided with ATM technology, which requires larger amounts of silicon to implement. Gigabit Ethernet will provide prioritized QoS (or CoS) based on emerging protocols such as IEEE 802.1Q/p and RSVP, which provide differentiated service levels, while ATM provides guaranteed QoS. ATM provides voice, data, and video integration, while Gigabit Ethernet delivers high-speed data. Voice and video capabilities over Gigabit Ethernet will depend on the success of video and voice over IP. Note that there is an applications overlap between Gigabit Ethernet and ATM. Both will be used for backbone, server, and building riser applications. However, only ATM can provide WAN services. Finally, ATM products are available today from Cisco. Fast EtherChannel is available today from Cisco as a migration path to Gigabit Ethernet for switch, router, and server connections. Gigabit Ethernet products ready for use in production networks will be available in 1998.

#### Where ATM Fits

Today, ATM is a robust technology for scaling campus networks, and products are available. ATM technology provides the following advantages for campus intranets:

- Guaranteed QoS
- Integration of data, video, and voice traffic (such as voice circuit emulation and data on the same backbone)
- WAN access/campus backbone integration to reduce operational cost
- Highest bandwidth available today on a single network interface (622 Mbps)
- Products are shipping

Cisco ATM LAN switches are being used today for campus backbone, building backbone, and server connections. Cisco customers cite the following reasons for deployment:

- Deployment of a scalable technology today
- Data-video-voice traffic integration
- WAN service integration
- Guaranteed QoS

For customers who are using ATM services in the WAN, this provides an opportunity to lower WAN costs through consolidation of ATM traffic at the WAN access point and across the campus backbone.

#### Where Gigabit Ethernet Fits

Gigabit Ethernet is in development with availability for production networks in 1998. Advantages offered by Gigabit Ethernet include the following:

- Low-cost bandwidth
- CoS based on the RSVP and the emerging IEEE 802.1Q/p standard, which provide differentiated service levels
- Leverage installed base of Ethernet, Fast Ethernet, and LAN protocols
- Leverage installed base of Ethernet knowledge for management, monitoring, and troubleshooting

Cisco Gigabit Ethernet products are designed for campus backbone, building backbone, and server connection solutions. For customers who will be deploying video and voice over IP, Gigabit Ethernet, along with Fast Ethernet, will provide a low-latency IP infrastructure to deliver these services throughout the enterprise campus. Until Gigabit Ethernet products are available, Fast EtherChannel can be deployed to scale bandwidth beyond Fast Ethernet for switch, router, and server connections.

#### Cisco Meets Industry Challenges with Complete Solutions


##### Cisco is Uniquely Qualified to Meet the Gigabit Networking Challenges

Cisco has been a pioneer and is now the industry leader in routing and switching solutions. In 1993, Cisco recognized that the convergence of routing and switching technologies is required for the delivery of high-performance, scalable campus networks. Today, Cisco is fusing router and switch technology to develop solutions that address the challenges of gigabit networking while providing a smooth migration path to gigabit network performance. This class of product is generally known as gigabit multilayer switching.

Cisco's gigabit networking initiative spans multiple Cisco groups, including Enterprise (which includes the former Granite Systems team), Internet Service Provider (ISP), and Small/Medium (SMB) lines of business. Each group addresses these challenges by leveraging Cisco technologies for delivery of multiple Gigabit Ethernet products with unparalleled performance for scaling campus networks to gigabit rates.

The Cisco Catalyst<sup>®</sup> LAN switching architecture provides the throughput and feature set required to scale intranet performance while it enables a smooth and stable migration of the network core, backbone, and server connections. Cisco has developed specific mechanisms to scale link layer, Layer 2, and network-layer performance. For gigabit networking, Cisco is evolving the Catalyst LAN switch products to increase system throughput up to 100 Gbps while providing full routing functionality with tens of millions of packets per second throughput, meeting the requirement to scale campus networks from Fiber Distributed Data Interface (FDDI), Fast Ethernet, or ATM to Gigabit Ethernet via a smooth migration path.

Figure 4 Catalyst Switch Evolution



Cisco is addressing the challenges of gigabit networking through the application of leading technology and product implementation:

- Multigigabit system bandwidth—Cisco's application-specific integrated circuit (ASIC)-based Catalyst LAN switch architecture, scalable to more than 100 Gbps of switching capacity, delivers Layer 2 switching performance with the option of line rate Layer 3 Forwarding—Layer 3 forwarding at Layer 2 performance. This feature enables users to select the type of throughput performance they need today and migrate performance and functionality as needed in the future.

- Network-layer (Layer 3) forwarding and routing at gigabit line rates— Cisco's NetFlow<sup>™</sup> LAN Switching, which incorporates route processing, ASIC-based forwarding, and Cisco IOS<sup>™</sup> software technologies, delivers line rate gigabit performance for Layer 3 forwarding and routing of strategic protocols such as TCP/IP while not compromising support for multiprotocol packet forwarding. NetFlow LAN Switching is based on standards and seamlessly integrates into current networks, providing network-proven scalable routing performance with industry-standard protocols.
- Application of network services at gigabit rates—Cisco's proven route processing and ASIC technology deliver the required network services at gigabit rates. Network services such as security, QoS, multimedia support, mobility, and policy implementation are enabled at gigabit speeds through deployment of NetFlow LAN Switching technology supported in Cisco's Catalyst switch products.
- Monitoring and management of gigabit systems—Cisco's innovative RMON implementation for the Catalyst LAN switch family is being extended to deliver monitoring, troubleshooting, and management for multigigabit systems. Unique application of RMON I and II will deliver tools for simple, effective resolution of gigabit networking problems.
- Smooth, scalable migration—Cisco's Catalyst LAN switch evolution assures a smooth migration to gigabit networking capability. This evolution will enable customers to migrate from FDDI, Fast Ethernet, and ATM to Gigabit Ethernet. Technologies such as Fast EtherChannel, Gigabit Ethernet, Gigabit EtherChannel and NetFlow LAN Switching will deliver performance improvements in manageable increments and provide the foundation for multigigabit networking performance. Cisco's current and future industry-leading Layer 2 and Layer 3 switching solutions ensure robust coexistence and interoperability with ATM solutions that span campus, metropolitan, and wide-area networks.

Table 4 Cisco Meets the Requirements

Requirement	Cisco Solution
Multigigabit switching system	Catalyst LAN switching scalable to 100 Gbps
Network-layer forwarding and routing at gigabit speeds	NetFlow LAN Switching, Cisco IOS multigigabit route processing
Application of network services at gigabit speeds	NetFlow LAN Switching, Cisco IOS software
Monitoring/management for multigigabit systems	Innovative RMON implementation, integrated tools
Smooth migration in manageable increments	NetFlow LAN Switching, Fast EtherChannel, Gigabit Ethernet, Gigabit EtherChannel technology
Uplinks, switches, router interfaces, gigabit switch routing	Catalyst LAN switches, Cisco 7500 router series, Cisco 12000 router series
Proven technology, scalable increments, low risk adoption	Proven Catalyst, Cisco IOS software, and route processing technologies

In summary, Cisco technologies including the Catalyst LAN switch architecture, NetFlow LAN Switching, Fast EtherChannel technology, Gigabit Ethernet, Gigabit EtherChannel technology, route processing, and Cisco IOS software will be combined to deliver gigabit-rate forwarding (Layer 2 and Layer 3), routing, and application of network services at gigabit rates to meet emerging application requirements. The ASIC-based Catalyst switching architecture will be scaled from current speeds to 100 Gbps. Fast EtherChannel technology, Gigabit Ethernet, and Gigabit EtherChannel technology deliver scalable link speeds between switching systems. NetFlow LAN Switching accelerates Layer 3 forwarding to line rate while it preserves the ability to apply network services. Cisco IOS software provides industry-leading multiprotocol routing and the rich feature set found in Cisco's proven routing systems. This unique combination of technologies ensures seamless migration from existing network designs to gigabit networks. Finally, this set of technologies will be implemented in the family of Catalyst switches. When combined with Cisco's award-winning product designs, a leading solution set is delivered for scaling campus intranets.

## Cisco Delivers Complete Gigabit Ethernet Solutions

Cisco will deliver Gigabit Ethernet products consistent with the completion of the IEEE 802.3z standard in 1998. Several strategic investments have been made in preparation for their delivery. These investments include:

- Acquisition of Granite Systems (provided Gigabit Ethernet MAC, low-cost ASIC, and gigabit switching technologies),
- Development of Gigabit EtherChannel technology
- Development of NetFlow LAN Switching technology
- Implementation of NetFlow LAN Switching in ASICs
- Development of gigabit Route Switch Processor (RSP) technology


Today, several product initiatives are underway to deliver a solution for scaling campus intranets to multigigabit performance. These initiatives include:

- Gigabit Ethernet uplinks for the Catalyst 5000 LAN switch family
- Gigabit Ethernet switching modules for the Catalyst 5000 LAN switch family
- High-performance route switch modules (RSMs) for the Catalyst family
- Gigabit Ethernet interfaces for the Cisco 7500 and 12000 series router product families
- Catalyst product enhancements that enable the scaling of both NetFlow LAN Switching and Layer 2 switching performance up to 100 Gbps
- Smooth migration for scaling performance on all high-end Cisco switching and routing product platforms

The latter initiative is particularly significant and represents an additional Cisco industry exclusive. In order to best allow users to effectively scale up performance in incremental, consumable steps, Cisco will leverage its existing Fast EtherChannel technology to provide its customers with a smooth evolution to gigabit networking and beyond. The combination of Cisco's Fast EtherChannel technology and Gigabit Ethernet will provide the foundation for scaling campus intranet bandwidth. Gigabit EtherChannel extends this capability to multigigabit speeds.

Cisco Fast EtherChannel technology builds on standards-based 802.3 full-duplex Fast Ethernet to provide reliable, high-performance solutions for the campus network backbone. Currently, Fast EtherChannel provides this bandwidth scalability by providing capacity increments from 200 Mbps to 800 Mbps. Cisco will extend these services to multigigabit capacity through Gigabit EtherChannel.

Figure 5 Cisco Ethernet Migration Paths



To complete the set of Gigabit Ethernet solutions, Cisco is working with leading industry partners to license and deploy Gigabit Ethernet interface technology for use in server adapters and test equipment. This low-cost, high-integration interface technology has resulted in a single-chip network interface card (NIC) design that excels in performance. Industry adoption of Cisco technology ensures high performance, interoperable server connections for gigabit networking, and early availability of test equipment for analysis and troubleshooting of gigabit networks.

Together, this set of Cisco and partner products form a complete solution for gigabit networking. The Cisco product set enables scaling of backbone performance based on network-layer functionality, which today enables thousands of enterprises to implement resilient, mission-critical networks for daily business operation on a global scale. Network-layer functionality is the critical element that delivers resilience, alternate paths, and network services that support client/server applications today and in the future.

As campus intranet performance requirements increase, Cisco products with accelerated network-layer performance will meet the challenge.

Ultimately, Cisco intends to distribute network Layer 3 functionality throughout the network core to ensure scalability and to enable a smooth migration path while taking advantage of the latest technology. For example, today, Cisco products employ Fast EtherChannel technology to scale link speed between systems, which leverages the investment in current switches, routers, and servers. Next, as the speed of the intranets increase and traffic patterns become unpredictable, route processing, Layer 3 forwarding, and services are distributed around the network to improve any-to-any communications and to leverage existing technology. This step is realized with NetFlow LAN Switching through use of the Catalyst RSM and NetFlow feature card. Finally, with further technology improvements in route processing, Layer 3 forwarding, and network service application, a seamless migration path is created toward fully distributed Layer 3 functionality at gigabit speeds. Here, Gigabit Ethernet serves as the equipment “interface” among the Layer 3 switching elements that are located within buildings and across the campus.

### Cisco Scales Technology Elements Required for Gigabit Networking

Three primary technologies must be scaled to gigabit capabilities for mass deployment of gigabit networking:

- Layer 3 functions—Layer 3 forwarding, route processing, and application of network services
- Layer 2 switching—Capacity or throughput available within the Layer 2 switching system
- Link bandwidth—Bandwidth available on the physical links between switches, routers, and servers

These technology elements must be scaled to enable product implementation that provides a smooth migration path, leverages installed equipment to extend useful life, and provides a clear path toward multigigabit network performance. The Layer 3 functions can be scaled in multiple ways that provide implementation options: one option is on a system basis; the second is by distributing Layer 3 functionality around the network. Link bandwidth and Layer 2 switching is scaled on a system-by-system basis.

Figure 6 summarizes the technology elements that must be scaled to deliver realistic migration steps that meet the demands of gigabit networking.

Cisco will scale each of these technologies and provide a smooth migration path via the Catalyst LAN switching family, Cisco 7500 router series, and the Cisco 12000 router series. NetFlow LAN Switching is Cisco's solution to scale Layer 3 forwarding and service application through flow caching. Key elements of NetFlow LAN Switching include Layer 3 forwarding (NetFlow forwarding), applications of network services (NetFlow services), and management (NetFlow management). The following sections summarize how Cisco will scale and deliver the technology elements required for gigabit networks.

**NetFlow LAN Switching Delivers Scalable Layer 3 Forwarding and Services**

Due to the rising levels of anywhere-to-everywhere communication, Layer 3 switching that scales to tens of millions of packets per second has become an imperative, required to speed up both peer-to-peer and client/server performance on campus networks. In addition to wire-speed performance, Layer 3 switching, must provide anywhere-to-everywhere connections without compromising latency, and must meet these additional requirements for scaling the campus:


- Transparently drop in and work in the LAN switching infrastructure

- Enable deployment in the LAN switched infrastructure with simple “plug-and-play” operation
- Preserve existing subnet structures
- Preserve the resilience, security, and scalability of traditional routed networks
- With Gigabit Ethernet on the horizon, scale to Gigabit speeds

Cisco responds to these challenges with NetFlow LAN Switching, which operates by switching flows (IP conversations) at Layer 3. The first packet of a flow takes the normal forwarding path. Information from the first packet is used to build an entry for “flow forwarding.” Subsequent packets are then Layer 3 switched at Layer 2 performance levels. In addition, NetFlow LAN Switching capitalizes on the flow nature of traffic to provide security and detailed statistic measurements. By distributing NetFlow LAN Switching technology onto the switches, Cisco increases the aggregate Layer 3 switching capacity of the campus intranet by an order of magnitude over current levels.

NetFlow Switching operates by creating a cache entry within a router, which contains the information needed to switch and perform appropriate services for each active flow. A flow is defined as a sequence of packets sent from a particular source to a particular destination. These packets are related in terms of their routing and any local handling policy they may require. After the NetFlow Switching cache is created, packets that are identified as belonging to an identified flow can be switched, based on the cached information, and appropriate services can be applied.

Figure 6 Cisco Scales the Technology for Evolution Toward Gigabit Networking




Extending the concept of NetFlow Switching beyond Cisco's routers, NetFlow LAN Switching now provides network-layer switching in Cisco's Catalyst series multilayer LAN switches at previously unmatched forwarding rates. This switching meets the bandwidth demands of tomorrow's next-generation backbone technologies. With development of NetFlow LAN Switching elements such as the Route Switch Module (RSM) and NetFlow Feature Cards, Catalyst 5000 series switches deliver millions of packet per second throughput performance. NetFlow LAN Switching can be deployed at any location in the network as an extension to existing routing infrastructures—from the campus backbone to the wiring closet. With NetFlow LAN Switching, network users can extend their use of Cisco IOS network services without paying the performance penalty usually associated with such processing-intensive functions. This increase in performance allows Cisco IOS network services to be utilized from end to end within the network and on a larger scale.

With NetFlow LAN Switching, the NetFlow feature cards silicon implements the line speed forwarding based on cached Layer 3 information. The RSM and Cisco IOS software continue to perform route processing functions on the first packet in each flow, while the Catalyst NetFlow feature cards forwards the subsequent packets in a flow (see Figure 7) at wire speed. This wire-speed forwarding architecture allows subsequent packets to bypass the route processor, which free route processor bandwidth, and preserve the value-added features and operating standards of Cisco IOS software.

NetFlow LAN Switching over LAN backbones uses multilayer switching silicon in the Catalyst series of LAN switches to automatically detect flows as they are switched within the Catalyst system, and it establishes a cut-through path whenever a flow is detected. Cisco IOS features of Catalyst LAN switches enable the switch to discover router information that is required for detecting candidate "flows," which, in turn, enable a RSM to inform the switch of routing or policy changes. Cisco IOS router software signals the NetFlow Feature Card to purge stale forwarding entries, for example, when an access list changes, and it enables switches to respond to network failures or topology changes to provide fast convergence and high availability.

Figure 7 NetFlow LAN Switching Operation




Because the implementation of NetFlow LAN Switching in Catalyst switches is based on Cisco ASIC technology, the Layer 3 throughput can be scaled from a few millions of packets per second to tens of millions of packets per second. When combined with the RSM processing power, based on high-end Cisco RSP technology, and Cisco's next-generation gigabit route processing engines, a complete solution for multigigabit Layer 3 forwarding and application of network services is delivered.

Cisco's solution to scaling Layer 3 forwarding and services employs multiple approaches. These approaches provide network architects a variety of options to achieve a smooth migration to multigigabit performance. The key elements of Cisco's NetFlow LAN Switching architecture for scaling are: route processing, multilayer switching, and NetFlow silicon (ASICs) implemented on the NetFlow Feature Cards. Each of these elements can be scaled within a system or around the network. Figure 8 illustrates these elements and the migration path for increasing Layer 3 functionality in gigabit networks. First, the Catalyst RSM delivers route processing and Layer 3 forwarding for deployment in Catalyst multilayer LAN switches. This feature enables the RSM to be placed at locations in the network that require increased Layer 3 functionality for scaling performance. When more throughput is required, increase the number of RSMs in the Catalyst switch.

Next, NetFlow ASICs are deployed via the NetFlow Feature Cards with the RSM and provide accelerated Layer 3 forwarding, millions of packets per second, at campus network locations that require high-performance Layer 3 switching. As route processing performance and ASIC technology increases these improvements will be incorporated into Catalyst LAN switches to deliver tens of millions of packets-per-second performance.

Figure 8 Scalable Layer 3 Forwarding and Services



### Cisco—The NetFlow LAN Switching Advantage

NetFlow LAN Switching goes beyond the basic customer criteria for Layer 3 switching and provides much more. It offers the following unique benefits:

- **Line rate performance**—With NetFlow LAN switching, Cisco addresses the need for wire speed Layer 3 Switching that can scale to gigabit speeds. It does so by embedding the Layer 3 switching in silicon.
- **Scalable**—NetFlow LAN Switching enables increased scalability in switched networks and allows customers to overcome the convergence and scalability problems of Layer 2 networks. By deploying NetFlow LAN Switching over a foundation of virtual LAN (VLAN) technologies, Cisco provides customers the benefits of both switching and routing on the same platforms.
- **Simple**—NetFlow LAN Switching preserves the simplicity and the cost effectiveness of the LAN switch infrastructure. It is autoconfigurable and autonomously sets up its Layer 3 flow cache. By its “plug and play” pragmatism, Cisco eliminates the need for users to learn new IP switching technologies.
- **Transparency**—It requires no end-system changes; no renumbering of subnets. It works with Dynamic Host Configuration Protocol (DHCP) and it requires no new routing protocols.
- **Standards based**—NetFlow LAN Switching is standards based. It uses Internet Engineering Task Force (IETF) standard routing protocols such as Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) for route determination. By being standards based, it can be deployed in a multivendor network.

- **Multiprotocol**—NetFlow Switching over LANs is based on the network-layer protocol standards such as IP, and it enables increased Layer 3 switching performance of strategic protocols such as IP. NetFlow LAN Switching supports multiprotocol environments and the continued use of all industry-standard, field-proven routing protocols available with Cisco IOS software.
- **Investment protection**—NetFlow LAN Switching will be available as a simple feature card upgrade on the Catalyst switches. Cisco enables customers to take advantage of NetFlow technology with the existing chassis and line cards on the Catalyst switch. In addition, it can use either an integrated RSM or an external router for route processing and Cisco IOS services. With NetFlow switching, Cisco demonstrates its commitment to preserving customers' investment in Cisco infrastructure.
- **Seamless migration**—NetFlow LAN Switching offers seamless migration paths to network managers and can transparently integrate into an existing network of Catalyst switches. It can be deployed incrementally in strategically located switches at the distribution or core of the campus network where performance improvements are required.

**Intelligent Cisco IOS Services with NetFlow LAN Switching**  
 In the NetFlow LAN Switching architecture, the RSM performs route processing and distributes Cisco IOS services to all ports on the Catalyst LAN switch. A lightweight management protocol run between the RSMs, the NetFlow Feature Card and the switch fabric enables the following services to be distributed:

- **Fast convergence**—In the campus core, the ability to recover from route failures and to Layer 3 switch around them is of paramount importance. NetFlow LAN Switching addresses this need for immediate convergence on routing topology changes. The performance gains delivered by NetFlow LAN Switching enable high-speed convergence required by gigabit networks. Route processors inform the switching fabric via a notification message of routing topology changes. The switch fabric responds to this notification by performing hardware-assisted invalidation of flow-forwarding cache entries.

- **Resilience**—NetFlow LAN Switching extends the benefits of Cisco's Hot Standby Router Protocol (HSRP) to the NetFlow LAN Switching without any additional configuration. This resilience enables the switches to transparently switch over to the hot standby backup router instantly when the primary router goes off line, and it eliminates a single point of failure in the network.
- **Access control**—An important attribute of classical routing in a network is the ability to set up access lists to filter or to altogether prevent traffic between members of different subnets. With the introduction of Layer 3 cut-through technology, the breaching of access lists and firewalls would seem likely. However, this is not the case. NetFlow LAN Switching is unique in enforcing multiple levels of security on every packet of the flow at wire speed. It allows the users to configure and enforce access control rules on the centralized router. Because NetFlow LAN Switching parses the packet up to the transport layer, it enables access lists to be validated. By providing multiple levels of security, it enables users to set up rules and control traffic based on IP addresses as well as transport layer application port numbers.
- **Accounting and traffic management**—A key requirement when deploying Layer 3 switching is to provide visibility of Layer 3 and Layer 4 flows as they are switched on the cut-through paths for troubleshooting, traffic management, and accounting purposes. NetFlow extends the benefits with management to address these needs. NetFlow enables detailed data collection of flow statistics maintained in hardware with no impact on switching performance. Flow information is made available to Cisco's Netsys Technologies™ applications, for example for network planning, RMON traffic management and monitoring, as well as to accounting applications.

#### **Multigigabit Layer 2 Switching Delivers Bandwidth for Network Segments**


Cisco's Catalyst LAN switch architecture scales to meet the challenge of gigabit networking. Today's award-winning Catalyst architecture delivers highly resilient Layer 2 switching through a combination of buffering, forwarding ASICs, and the switch architecture's unique implementation. The architecture scales today in port density, bandwidth, and chassis size and maintains a consistent set of features. This Layer 2 architecture is delivered today by the Catalyst 5000, 5002, and 5500 switches.

Scaling Layer 2 switching performance is useful to a point. Eventually, Layer 3 forwarding performance must also be improved. Cisco's architecture provides the flexibility to scale both Layer 2 forwarding and Layer 3 forwarding performance at the same time. The NetFlow LAN Switching architecture enables ASIC-based forwarding on Layer 2 or Layer 3 information. As the forwarding requirements of networks extend beyond the performance that is possible with processor-based Layer 3 forwarding, the integration of Layer 2 forwarding and Layer 3 forwarding in silicon becomes a critical performance element. Here, Layer 2 provides the necessary multiprotocol-based switching while the Layer 3 forwarding is focused on strategic network protocols such as IP to gain performance advantages. Layer 3 provides fast convergence, alternate path routing, performance increases with system integration, and application of network services at high speed.

The Catalyst LAN switch's multilayer architecture for (Layer 2 forwarding and Layer 3 forwarding) scales to 100 Gbps. Cisco has designed the Layer 2 forwarding and Layer 3 forwarding elements of this architecture to scale with ASIC processes and integration. The benefit is a consistent set of features and network services based on world-class ASIC technology that assures low cost and performance migration. The scaling steps provide a smooth migration path that enables network managers to choose the best solution for their situation to achieve low cost of ownership while maintaining flexibility for future growth.

Cisco's multilayer architecture implementation is specifically designed for combining Layer 2 switching with Layer 3 switching to deliver accelerated performance of Layer 3 forwarding, while it continues to deliver network services. This Layer 3 functionality is delivered via NetFlow LAN Switching through the RSM and NetFlow feature card incorporating ASIC-based forwarding. As previously discussed, Layer 3 functionality is an important aspect of scaling campus networks. Migration to Layer 3 switching functionality is an important benefit of Catalyst LAN switch deployment. The multilayer LAN switch architecture allows Layer 3 forwarding and services to be deployed where required to scale the campus network.

Figure 9 Scaling/Integration of Layer 2 and Layer 3 Switching




#### Fast EtherChannel Delivers Link Scalability—Migration to Gigabit Ethernet and Beyond

Scaling “links” between equipment to gigabit networking speeds is accomplished on a system-by-system basis. Use of industry-standard interfaces maximizes flexibility. Multiple links leverage existing and future technology and standards. Cisco's Fast EtherChannel technology delivers link scalability based on industry-standard interfaces.

Cisco is committed to providing its customers with smooth evolution to Gigabit Ethernet bandwidth and beyond. The combination of Cisco's Fast EtherChannel technology and Gigabit Ethernet will provide the foundation for scaling campus intranet bandwidth. Cisco's Fast EtherChannel technology builds upon standards-based 802.3 full-duplex Fast Ethernet to provide reliable high-speed solutions for the campus network backbone. Fast EtherChannel technology provides bandwidth scalability within the campus by providing increments from 200 Mbps to 800 Mbps with multigigabit capacity in the future. Fast EtherChannel technology not only solves the immediate problem of scaling bandwidth within the network backbone today, it also paves the path for an evolution to standards-based Gigabit Ethernet and beyond, because Fast EtherChannel technology can be applied to support Gigabit EtherChannel technology.

Fast EtherChannel technology provides a solution to higher bandwidth between servers, routers, and switches than Fast Ethernet technology can currently provide, and it will extend to deliver Gigabit EtherChannel technology that provides multigigabit backbone links in the future.

Figure 10 Scalable Link Bandwidth with Cisco Fast EtherChannel Technology



Today, Fast EtherChannel technology provides scalable, incremental bandwidth that you will not have to replace after Gigabit Ethernet standards are finalized. Fast EtherChannel technology provides the following benefits:

- **Standards based**—Fast EtherChannel builds upon IEEE 802.3u-compliant Fast Ethernet by grouping multiple full-duplex point-to-point links together (up to four). Fast EtherChannel technology uses the industry-standard 802.3x mechanisms for full-duplex autonegotiation and autosensing.
- **Multiple platforms**—Fast EtherChannel technology is flexible and can be used anywhere in the network where bottlenecks are likely to occur. It can be leveraged in network designs to increase bandwidth between switches, between routers and switches, and can provide scalable bandwidth into network servers such as large UNIX servers or Windows NT-based Web servers.
- **Flexible incremental bandwidth**—Fast EtherChannel technology provides bandwidth aggregation in multiples of 200 Mbps, with multiple Gigabit Ethernet interfaces in the future. For example, network managers can deploy Fast EtherChannel technology that consists of pairs of full-duplex Fast Ethernet to provide 400+ Mbps between the wiring closet and the data center. Meanwhile, in the data center, bandwidths of up to 800 Mbps can be provided between servers and the network backbone to provide large amounts of scalable, incremental bandwidth.
- **Load balancing**—Fast EtherChannel technology is composed of multiple Fast Ethernet links and is capable of load-balancing traffic across those links. Unicast, broadcast, and multicast traffic is evenly distributed across the links that provide higher performance and redundant parallel paths. In the event of a link failure, traffic is redirected to remaining links within the channel without user intervention.

- **Resiliency and fast convergence**—Fast EtherChannel technology provides automatic recovery for loss of a link by redistributing loads across remaining links. If a link does fail, Fast EtherChannel technology redirects traffic from the failed link to the remaining links within less than a second. This convergence is transparent to the end user—no host protocol timers expire; therefore, no sessions are dropped.
- **Ease of management**—Fast EtherChannel technology leverages the experience and know-how developed over the years in troubleshooting and maintaining Ethernet and Fast Ethernet networks. Existing network probes can be used for traffic management and troubleshooting, and existing management applications such as CiscoView and TrafficDirector™ will be leveraged to support Fast EtherChannel technology.
- **Transparent to network applications**—Fast EtherChannel technology does not require any changes to networked applications. When Fast EtherChannel technology is used within the campus, switches and routers provide load balancing across multiple links transparently to network users. For support of Fast EtherChannel technology on enterprise-class servers and NIC cards, smart software drivers can coordinate distribution of loads across multiple network interfaces.
- **Embedded in Cisco IOS software**—Fast EtherChannel technology connections fully utilize Cisco IOS VLAN and routing technologies. Inter-Switch Link (ISL) VLAN trunking protocol can carry multiple VLANs across a Fast EtherChannel device, and routers attached to Fast EtherChannel trunks can provide full multiprotocol routing with support for hot standby using HSRP.
- **Gigabit Ethernet ready**—Fast EtherChannel technology is Gigabit Ethernet ready. Fast EtherChannel technology allows network managers to deploy networks today that scale smoothly to the availability of standards-based Gigabit Ethernet and beyond with future Gigabit EtherChannel technology.

## Campus Networks Designed to Scale

Network managers are confronted with challenges in migrating their networks to meet tomorrow's requirements. Cisco provides a blueprint to smoothly evolve with a multilayer network architecture and meet the changing requirements of campus networks. Network managers increasingly need to deploy scalable, flexible networks that will accommodate growing demands for bandwidth, stability, and manageability, while minimizing the financial impact on their existing network infrastructures.

High-performance network elements are deployed where needed to provide the performance required to meet the business need. Driven by these requirements, the network industry has evolved toward a new network architecture, the switched internetwork.


Table 5 Campus Network Trends and Network Design Implications

Trends	Network Design Implication
Any-to-any traffic	L2 delivers limited scaling, L3 scales
Increasing network size	L2 delivers limited scaling, L3 scales
Increasing throughput performance	L3 = L2 performance
Network services	L3 network services application
Network resilience	L3 for fast convergence and redundancy
Simple migration steps	Network evolution over time

The design of campus networks today must address the scaling challenges faced by today's networks. These include any-to-any traffic, increasing network sizes, increasing throughput performance demands, requirements for network services, resilience and simple migration steps. Each of these has specific implications on the network design as summarized in Table 5 above. Networks have evolved from router-based backbones with shared-media hubs in the wiring closets to LAN switch and router backbones with LAN switches in the wiring closets. At each step along the way, Layer 3 or routing has provided important network functionality, including: establishment and management of the network hierarchy and topology, controlled forwarding based on network-layer address for multicast and QoS, controlled access based on network-layer address, and troubleshooting. Layer 3 functionality has enabled scalability of the Internet and campus enterprise networks today. Nearly all of today's trends lead to the addition of Layer 3 functionality to improve and scale performance of campus networks. Network managers must evolve their networks to meet requirements over time.

By implementing campus network designs with scaling as the primary objective, network managers are well positioned to deliver performance and services required for meeting the challenge of geometric traffic growth. Cisco has mapped out a series of steps which provide easy migration for performance enhancement as enterprise networks evolve toward the intranet model. These steps are summarized in the following sections.

Figure 11 Layer 2 End-to-End VLANs with Ethernet Backbone



### Networks Today


Many of today's frame-based networks employ Fast Ethernet and end-to-end VLAN technology. In this design users are grouped into VLANs (subnets) independent of physical location. All users in a VLAN have the same traffic pattern, 80 percent local and 20 percent to other VLANs (subnets). As a person moves around the campus, their VLAN membership follows them and remains the same. Each VLAN has a common set of security requirements for all members. Communication between VLANs or user groups is provided by inter-VLAN routing. Note that all VLANs cross all switches for end-to-end connectivity. This is illustrated in Figure 11. While this is a cost-effective implementation, as traffic patterns change toward any-to-any, more users are added and the traffic grows geometrically, the ability to scale this design is limited until Layer 3 is introduced.

As the network size and traffic patterns change, a Layer 2/Layer 3 network must be introduced into the distribution layer to meet the traffic demands. This is the first step in getting the network ready to scale for evolution to gigabit performance. Here the network equipment at the core and distribution layer are connected through Fast Ethernet. Layer 3 throughput is accelerated with NetFlow LAN Switching technology via deployment of the RSM in Catalyst switches. This design accommodates the new traffic patterns, which are a departure from the 80 percent local and 20 percent over the backbone. In this case, the user's VLAN changes with physical moves. DHCP is used to provide client mobility throughout the network. Most enterprise servers are located centrally at a few locations within the enterprise. High performance is delivered by distribution of Layer 3 functionality around the network. This is illustrated in Figure 12.

Intranet servers have become an important tool in disseminating information from one workgroup to many, causing the phenomenon of any-to-any traffic flows. The network architecture must evolve to handle this new trend. As more intranet servers are deployed, more traffic flows between VLANs. Whereas earlier 80 percent of the traffic remained within the VLAN, now only 60 percent or less remains in the VLAN, 40 percent now flows across VLANs. To scale to these new traffic patterns, networks need more inter-VLAN forwarding power, as well as more bandwidth in the backbone. To solve the bandwidth challenge, Cisco's Fast EtherChannel technology provides up to four times the bandwidth of Fast Ethernet, yet works with industry standards. To support the increased inter-VLAN traffic flows and maintain wire-speed performance, Cisco delivers multiple millions of packets per second performance with NetFlow LAN Switching via the RSM and NetFlow feature card for Cisco's Catalyst switches.

Implementation of the following three phases will depend greatly on the throughput and bandwidth needs of the specific campus network. Each step provides improved performance and can be employed as needed.

Figure 12 Layer 2/Layer 3 Hierarchy Design



**Phase 1—Improve Bandwidth and Layer 3 Throughput**  
Phase I improves bandwidth in the core and distribution layers of the network. The first scaling step is to increase bandwidth, leverage the investment of installed equipment, and employ proven technology. Fast EtherChannel technology is deployed in the core and distribution layers to increase throughput and fully utilize installed equipment. Fast EtherChannel is proven and shipping today. It is used to increase bandwidth where needed to meet growing requirements. This step is particularly attractive for customers who are waiting for Gigabit Ethernet technology to mature before deployment. This step is illustrated in Figure 13.


Widespread use of DHCP for dynamic IP addresses has accompanied the paradigm shift to intranet servers and any-to-any traffic flows. With DHCP-based assignment of IP addresses, clients are no longer guaranteed a direct Layer 2 connection to their servers. This scenario increases the need for faster Layer 3 interconnect while reducing the necessity of VLANs as a mobility solution. Improving Layer 3 performance is accomplished with the deployment of the NetFlow feature card.

**Phase 2—Gigabit Ethernet Switching in the Core**  
 Phase 2 is the addition of proven gigabit Layer 2 switching based on Gigabit Ethernet interfaces where needed in the network. These Layer 2 devices are placed in the core with links to the distribution layer switches and connected via Gigabit Ethernet uplinks. Core switches are connected via Gigabit EtherChannel delivering multigigabit bandwidth between systems. Enterprise servers are connected

with Fast EtherChannel, Gigabit Ethernet or Gigabit EtherChannel technology delivering high throughput for mission-critical intranet and enterprise applications. While this phase is effective for increasing bandwidth in the core, there is limited scalability of the network core due to the Layer 2 functionality. This is illustrated in Figure 13.

For high-speed connections from the distribution layer, to the wiring closet, Fast EtherChannel technology is the first step to increase the uplink bandwidth. If very high performance is required, Gigabit switching can be added to the distribution layer and Gigabit Ethernet can be used in the building risers. Here, Gigabit Ethernet links in the building data center are connected to Gigabit Ethernet uplinks in high-performance wiring closets.


Figure 13 Phase I—Improve Bandwidth and Layer 3 Throughput



Desktop CPU performance continues to increase, and more users are beginning to run at Fast Ethernet speeds, putting more pressure on the network to scale and support this increased growth in users, subnets, and bandwidth. As Gigabit Ethernet becomes standardized and available, performance of Layer 3 switching services needs to be increased. Faster-performing desktops and exploding intranet server applications continue to increase any-to-any traffic flows. These new applications further reduce the amount of traffic within a VLAN to 40 percent, pushing 60 percent to Layer 3 inter-VLAN flows. Also with the increase in intranet servers and applications comes the need for more traffic control and visibility into the flows.

Cisco addresses this need for wire-speed Layer 3 switching by use of NetFlow LAN Switching in the distribution layer. This provides high-performance IP forwarding along with Cisco IOS-based multiprotocol routing, access lists, and QoS services for control and flow management for visibility and planning.

Figure 14 Gigabit Ethernet Switching in the Core




### Phase 3—Gigabit Multilayer Switching in the Core

The ultimate goal is to deploy gigabit multilayer switching in the core for high-performance throughput for any-to-any traffic and long-term scalability. With gigabit line rate performance for Layer 3 forwarding and service application, this solution delivers the long-term scalability necessary for any-to-any traffic, geometric traffic growth, and large network implementations. Further increases in bandwidth demands will be addressed by the use of Gigabit EtherChannel technology for interconnection of switches and enterprise servers. This is illustrated in Figure 15.

By this time, intranet applications are wide spread with high volumes of traffic flowing over the backbone. An estimated 80 percent of the enterprise traffic will traverse the backbone, while 20 percent of the traffic remains in the local workgroups. This is driven by broad use of applications based on Web technology. A simple click of the mouse will summon files from enterprise servers at a variety of locations around the campus; thus, all files will traverse the backbone to reach the client. VLANs will be confined to use in wiring closets for grouping users with similar characteristics or policies, and the distribution layer for traffic load balancing and redundancy. Mobility will be delivered through enterprise deployment of DHCP.

Figure 15 Gigabit Multilayer Switching in the Core




## Cisco's Path to Gigabit Multilayer Switching for the Campus

### The Advent of Gigabit Multilayer Switching

The collection of requirements discussed thus far has been responsible for the definition of gigabit multilayer switching—a blend of the high performance delivered by hardware-based switching with the required control and value-added services that are provided by traditional Layer 3 routing. The delivery of gigabit multilayer switching eliminates the need to segregate switching and routing functionality to deliver high performance. With gigabit multilayer switching, network managers can finally scale campus network performance with confidence and control in the face of any-to-any traffic patterns and geometric traffic growth. They can now address high-performance application requirements without compromising network control, scalability, and speed.

An additionally enticing attribute of gigabit multilayer switching is its inherent technology independence. It does not depend strictly on frames or cells for proper operation. Either can be supported to most effectively enable what each does well—all in the context of a superior technical framework. For network managers, gigabit multilayer switching designed for the campus delivers a blend of gigabit throughput and routing functionality based on proven Cisco routing technology, which powers 80 percent of today's campus networks. Cisco has combined Gigabit Ethernet with technology advances in routing to deliver a solution set that meets the challenges of geometric traffic growth and any-to-any traffic patterns.

Figure 16 Gigabit Multilayer Switching Scales the Campus Intranet



While it is important to keep the goal of gigabit networking in mind, the migration path to get there is a key consideration, especially if there are budget constraints and plans to leverage installed systems. Cisco has outlined a migration path for products and network designs that evolve smoothly, easily leveraging installed equipment and providing realistic increments for managing budgets.

Table 6 Smooth Migration Steps Toward Gigabit Networking with Cisco

Today—Phase 1	Phase II	Phase III
<b>Goals:</b> <ul style="list-style-type: none"> <li>• Multi-100 Mbps link speeds</li> <li>• Increase network-layer performance</li> </ul>	<b>Goals:</b> <ul style="list-style-type: none"> <li>• Gigabit link speed</li> <li>• Millions of packets per second network-layer performance</li> </ul>	<b>Goals:</b> <ul style="list-style-type: none"> <li>• Multigigabit link speeds</li> <li>• Tens of millions of packets per second performance</li> </ul>
<b>Network Design:</b> <ul style="list-style-type: none"> <li>• Establish scalable architecture</li> <li>• Add Layer 3 to the distribution layer</li> </ul>	<b>Network Design:</b> <ul style="list-style-type: none"> <li>• Extend scalable architecture</li> <li>• Layer 3 performance to 2-3 Mpps</li> </ul>	<b>Network Design:</b> <ul style="list-style-type: none"> <li>• Layer 3 performance to tens of Mpps</li> <li>• Gigabit multilayer switching in the core</li> </ul>
<b>Solutions Elements:</b> <ul style="list-style-type: none"> <li>• Fast EtherChannel</li> <li>• NetFlow LAN Switching</li> <li>• Route Switch Module</li> </ul>	<b>Solutions Elements:</b> <ul style="list-style-type: none"> <li>• Fast EtherChannel</li> <li>• Gigabit Ethernet</li> <li>• Route Switch Module</li> <li>• NetFlow Feature Card</li> <li>• GE Switching</li> </ul>	<b>Solutions Elements:</b> <ul style="list-style-type: none"> <li>• Gigabit Ethernet</li> <li>• Gigabit EtherChannel</li> <li>• Route Switch Module</li> <li>• NetFlow Feature Card</li> <li>• Gigabit multilayer switch</li> </ul>

With Cisco's Catalyst products, network managers can take advantage of migration in incremental steps, as needed, to improve performance of campus intranets. New Catalyst product elements need only be deployed in the areas of the network that require performance improvements. Table 6 summarizes the goals, network design and the Cisco solution elements.

For long-term performance gains, Cisco's approach to the implementation of gigabit multilayer switching for campus intranets is called NetFlow LAN Switching. With NetFlow LAN Switching—routing, switching, and network-layer services are blended to deliver a compelling combination of high performance and high functionality.

Routing functions, for their part, look outward at the topology and hierarchy of the network to calculate routing tables, control forwarding, and enable network scalability with unpredictable traffic patterns. However, rather than having to then carry the entire burden of performing all traffic control functions, NetFlow LAN Switching allows routing decisions to take place and transmits the appropriate forwarding and service application information to silicon-resident (or ASIC-based) switching functions. These fast, intelligent ASICs can quickly refer to this information as required to execute high-speed packet forwarding—thus avoiding the need to send every packet through a route processor; yet providing full routing functionality to every packet in a flow. NetFlow LAN Switching is the best of both worlds. Cisco's implementation of NetFlow LAN Switching for the Catalyst LAN switch has two components today: the RSM and the NetFlow Feature Card. The RSM performs route processing and determines specific services to be applied based on the first packet in a flow. The resulting information is stored in the NetFlow cache on the NetFlow Feature Card. The remainder of the flow is then forwarded via the NetFlow Feature Card's ASIC-based forwarding engine with all applied network layer services. Both the RSM and the NetFlow Feature Card are based on field-proven Cisco technology. As these elements are scaled in performance to gigabit line rate performance, together, they will deliver full routing functionality to every packet at multigigabit speeds.

## Conclusion

We have discussed how today's unique networking environment is becoming increasingly driven by Internet and intranet services. Clearly, the result is the creation of accelerated network performance demands and unpredictable traffic patterns. More specifically, it is this multidimensional combination of application, system capacity, and user community demands that is driving annual capacity growth in many user networks to triple-digit rates. The resulting success metrics for network managers are defined by their ability to implement a networked infrastructure whose capacity can consistently remain ahead of user demands at any time.

Cisco believes that this collection of user-performance requirements is rapidly driving the industry toward a gigabit networking model. Unlike other vendors that define solutions to these performance requirements only in the context of one technology or product, Cisco is taking a much broader approach in its unique combination of end-to-end gigabit multilayer switching solution that leverages proven Cisco technology to achieve the best performance, functionality, service application, migration, and manageability.

As always, Cisco continues to provide customers with intelligent, high-performance systems that deliver value-added end-to-end services. Fast EtherChannel technology, Gigabit EtherChannel technology, and NetFlow LAN Switching are just three examples of how Cisco is leading the industry by offering a smooth, comprehensive migration to next-generation networks. The implication is clear: When considering how to best implement a gigabit networking solution for your business, come to the leader who defines the complete solution rather than just a piece of it. Come to Cisco first.



**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems Europe s.a.r.l.  
Parc Evolic, Batiment L1/L2  
16 Avenue du Quebec  
Villebon, BP 706  
91961 Courtaboeuf Cedex  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 6918 61 00  
Fax: 33 1 6928 83 26

**Americas**

**Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Headquarters**

Nihon Cisco Systems K.K.  
Fuji Building, 9th Floor  
3-2-3 Marunouchi  
Chiyoda-ku, Tokyo 100  
Japan  
<http://www.cisco.com>  
Tel: 81 3 5219 6250  
Fax: 81 3 5219 6001

**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at <http://www.cisco.com>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark  
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia  
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland • Singapore  
South Africa • Spain • Sweden • Switzerland • Taiwan, ROC • Thailand • Turkey • United Arab Emirates • United States • Venezuela