3Com®

3TECH™

The 3Com Technical Journal Volume 6 • Number 2 April 1995

Inside This Issue:

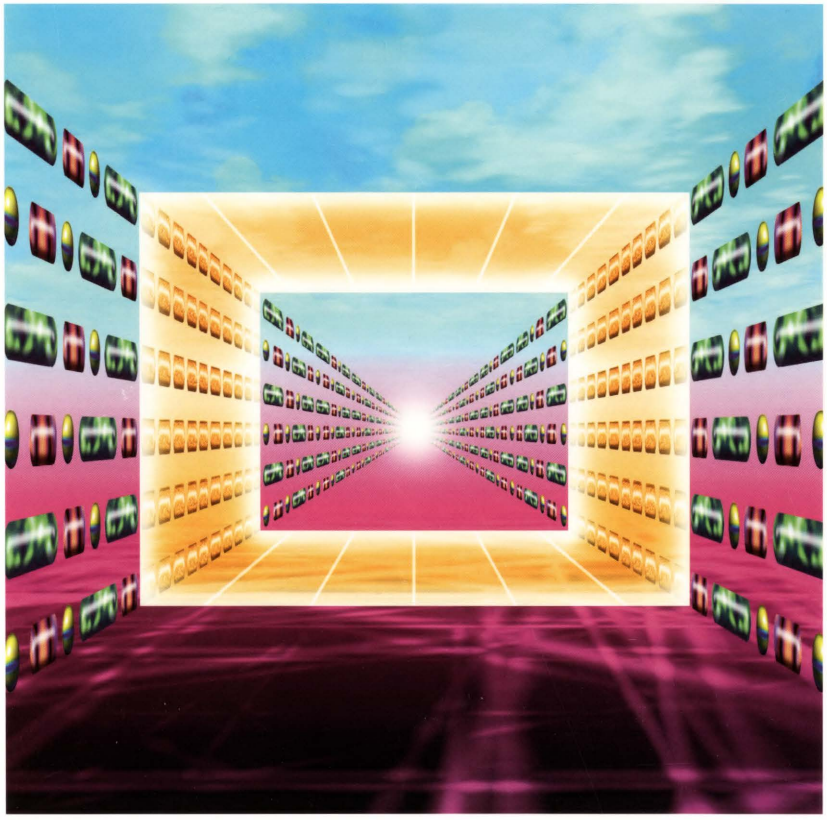

Integrating ATM Across the Enterprise Data Network

3TECH The 3Com Technical Journal

3TECH, 3Com's technical journal, is published quarterly by 3Com Corporation, Santa Clara, California 95052.

Officers: Eric A. Benhamou, president and chief executive officer; Bob Finocchio, executive vice president, network systems operations; Christopher B. Paisley, vice president and chief financial officer; Debra Engel, vice president, corporate services; John Hart, vice president and chief technical officer; Doug Spreng, vice president and general manager, network adapter division; Alan Kessler, vice president, systems sales–North America; Janice Roberts, vice president, central marketing; Ralph Godfrey, vice president, volume sales–Americas; Richard Joyce, vice president, sales–Asia Pacific Rim.

Circulation and Policy

3TECH is sent free to qualified subscribers. Send qualification card (enclosed in this issue) and address changes to: *3TECH*, P.O. Box 869038, Plano, Texas 75086-9870 or fax them to 408-764-6477. For customer service questions, call 408-764-6626.

Submissions

Manuscript submissions, inquiries, and all other correspondence should be addressed to *3TECH*'s editor: Suzanne Dowling, 3Com Corporation, P.O. Box 58145, Santa Clara, California 95052-8145. Articles in *3TECH* are primarily authored by 3Com employees; however, non-3Com authors are encouraged to submit for publication articles dealing with 3Com-related research or solutions to technical problems.

ISSN 1051 9637

Printed on recycled paper with soy-based inks

3TECH The 3Com Technical Journal

Volume 6, Number 2 • April 1995

Feature Articles

Departments

Inside:

Features

Departments

Announcements

Editor's Note

Do You Know Where Your Bottleneck Is?

By Jose Fraga, Guest Contributor

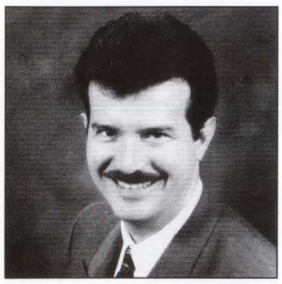

Jose Fraga is a network consultant in 3Com's Boca Raton, Florida office. Prior to joining 3Com, Jose was a network specialist for Microsoft's consulting services division.

He received his B.S. in electrical engineering from Universidad Metropolitana in Caracas, Venezuela.

Locating bottlenecks used to be a simple, straightforward task. Shared-media, single-protocol, single-server networks were the norm. Productivity applications invoked occasional file transfers between server and workstation. Network management tools could easily see all network traffic, and it was relatively simple to identify and correct performance problems—a new adapter, a more powerful server, or a faster PC usually did the trick.

But workgroups grew and networks grew. Groupware and client/server applications emerged, changing network traffic patterns and pushing LAN segments to much higher levels of utilization. As departmental LANs were interconnected, multiple transport protocol and media support further complicated network management. Bottlenecks became increasingly difficult to track down. All the while, users were unrelenting in their demands for higher performance.

Network managers have a vast array of choices for increasing performance. Segmenting shared-media LANs and centralizing file servers isolates client/server workgroup traffic and increases aggregate bandwidth. High-speed technologies such as FDDI and 100BASE-T offer enormous performance improvement on any segment. And now, as we incorporate connection-oriented technologies like ATM into our shared media LANs, we can virtually guarantee that the network itself is not the bottleneck.

Identifying the Real Problem

But before introducing any new technology to your network, it is important to correctly identify the source of the problem. For example, more network bandwidth won't solve a performance problem if the bottleneck is at the workstation level (see figure). We need to take a rigorous, system-wise approach to locating bottlenecks, and then choose the solution wisely.

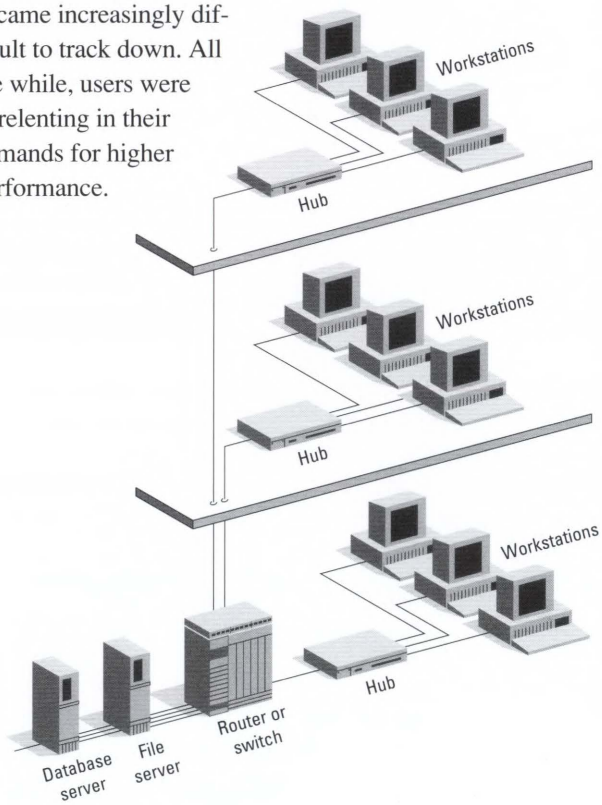

Potential Workstation Bottlenecks
LAN adapter
PC bus
CPU speed
Available RAM
Protocol stack
Operating system
Client application

Potential Network Bottlenecks
Segment bandwidth
Backbone bandwidth
Router speed
Multihop latency
Broadcast storms
Protocol inefficiencies

Potential Server Bottlenecks
LAN adapter
PC bus
CPU speed
Available RAM
Hard disk speed
Protocol stack
Operating system
Network operating system
Server applications
SQL engine

Isolating bottlenecks involves breaking the system down into its subsystems. For each subsystem, we must further understand the role of each of its components. Workstation performance, for example, is not only a function of CPU speed, but also of RAM size and speed, bus architecture, the mode of data transfer between adapter and host, transport protocol, operating system, and application software. End-to-end performance is the result of a complex chain of events, and performance is only as good as the weakest link in the chain.

There are many tools to help track down bottlenecks. Because every system is unique, no single tool can identify all problems. Network operating system tools provide the most insight into how server CPU cycles are being utilized. Tools like the Windows NT™ Performance Monitor help isolate workstation problems. And of course, network management software like 3Com's Transcend® pinpoints network-specific problems, and provides the best view of the system as a whole.

The Right Solutions

Once a bottleneck is correctly identified, it is important to distinguish theory from reality when evaluating potential solutions. For example, an FDDI adapter on an EISA bus can theoretically sustain throughput as fast as 64 Mbps. But specifications like this are often based on laboratory conditions; real-world throughput will almost certainly be less.

Furthermore, the very same hardware will respond differently in different software environments. For example, an adapter with a bus-master DMA data transfer scheme will minimize the load on a server's CPU, freeing it for other important tasks. But the same card may not improve performance in a workstation that simply idles while it waits for a server's response. In fact, if the card uses a large amount of host RAM in an already overcrowded DOS machine, performance may actually suffer.

Coping with Complexity

As system architectures evolve, powerful new tools are emerging for managing them. But systems will still be complex, and each installation is unique.

3Com has responded with powerful tools that allow proactive network management. With SmartAgent™ management agents embedded in 3Com routers and hubs, managers can customize Transcend to notify them about specific events in which they are interested. Alarms and events on all ports help network administrators isolate problems—even on a switched network where a sniffer might be more difficult to use effectively.

Complexity is also held in check by reducing the number of component vendors. By choosing as many network components as possible from a single vendor, you ensure better network integration and optimize your network's performance. For example, RMON network management with 3Com hubs, routers, and adapters provides a more complete management picture than a multivendor approach. And when you go with equipment from a single vendor, you get a consistent user interface and a single source for service and support, simplifying your network operations.

Of course, there is no substitute for knowing your system and staying informed about the ever-changing technology landscape. Feature articles in this issue of *3TECH* offer valuable insight into several key networking issues that can affect capacity and performance: ATM, intelligent switching, and SmartAgent gauges. But to maximize your resources, be sure to correctly identify your system's bottleneck, and then select the best solution for the job. ☐

Feature Article

Integrating ATM Across the Enterprise Data Network

ATM LAN Emulation As the Key First Step

By Bob Klessig

This article provides practical information for network managers and administrators planning to migrate to Asynchronous Transfer Mode (ATM) as a high-speed data networking technology for their growing corporate networks.

The article reviews basic ATM concepts, then describes the LAN Emulation standard defined by the ATM Forum, which allows existing applications to access an ATM network using traditional network protocol stacks. It explains how LAN Emulation enables the construction of virtual LANs for logical grouping of workgroups across the network. Finally, it explores two deployment scenarios for scaling ATM into the enterprise—first in campus and WAN backbones, and then in the building backbone, workgroups, and server farm.

ATM migration is a key component of 3Com's High-Performance Scalable Networking (HPSN) architecture. HPSN addresses the demand for higher bandwidth by delivering the speed of ATM while leveraging existing LAN technologies. 3Com's Transcend® family of network management solutions provides visibility in the physical and logical relationships of a growing enterprise network. By adopting an evolutionary strategy, network managers can integrate ATM smoothly across all functional areas of the network as their needs require and as ATM technologies mature.

Bob Klessig is manager, business development at 3Com Corporation. He is an expert in high-speed networking in the areas of FDDI and ATM. He is the principal company representative in the ATM Forum Technical Working Group with emphasis on LAN emulation, traffic management, and WAN services.

Prior to joining 3Com in 1992, Bob was district manager of data protocol architecture planning for Bell Communications Research. Prior to divestiture, he was a supervisor at Bell Telephone Laboratories.

Bob holds a Ph.D. and M.S. in electrical engineering and computer sciences from the University of California, Berkeley.

ATM Basics

As companies began interconnecting LANs across wide areas in the early 1980s, a dichotomy of networking technologies emerged: private local area networks based on shared media and optimized for data transmission, but connected over public telephony networks based on dedicated bandwidth and optimized for voice traffic. In 1986, the Comité Consultatif International Télégraphique et Téléphonique (CCITT), now known as the International Telecommunications Union (ITU), formed a study group to explore the concept of a high-speed, integrated network that could uniformly handle voice, data, and a variety of other services. The result of their deliberations is BISDN, or the Broadband Integrated Services Digital Network. BISDN services require high-speed channels for transmitting digitized voice, data, video, and multimedia traffic. Asynchronous Transfer Mode (ATM) is the switching and multiplexing technology for supporting BISDN services.

Accommodating Bursty and Real-Time Communications

One of the greatest challenges in defining ATM was to determine a structure that could efficiently handle any type of traffic. Such a structure must accommodate a variety of bit rates and support bursty communications, since voice, data, and video traffic all exhibit bursty behavior.

While most people assume that circuit-switched voice traffic is not bursty, the acoustic energy generated by one side of a voice conversation is present only about 40 percent of the time. Undersea telephone transmission systems have doubled their voice capacity for years by exploiting this fact and allowing each voice circuit to transmit only during active periods.

Packet-switching has been the technology of choice for bursty data traffic because it consumes bandwidth only when traffic is present. But traditional packet-switching mechanisms cannot achieve the performance and speed required for real-time, two-way traffic. ATM overcomes this limitation by offering fixed-length packets. Each ATM packet, called a cell, consists of a 48-byte payload and a 5-byte header. Fixed-length ATM cells offer several advantages:

- Networking and switching queuing delays are more predictable with fixed-length data cells. Switch vendors can put mechanisms in place to ensure the appropriate level of service for all types of traffic, especially for delay-sensitive services such as voice and video.
- It is less complex and more reliable to process ATM cells than variable-length packets. High predictability allows ATM hardware to be implemented more efficiently, because control structures, buffers, and buffer management schemes can be designed to known size criteria.
- Fixed-length cells allow cell-relay switches to process cells in parallel, for speeds that far exceed the limitations of bus-based switch architectures.

Like traditional packet data, ATM cells require bandwidth only when traffic is present, yet they can provide the equivalent of a time-division multiplexer time slot for continuous traffic like today's digitized voice. As a result, ATM can handle real-time and bursty LAN traffic equally well.

Efficient bandwidth use is not the only issue addressed by ATM technology. In fact, different traffic types require different delay behavior, delay variation, and loss characteristics. ATM provides different qualities of services to accommodate these differences. To access the network, a station requests a virtual circuit between the transmitting and receiving ends. During connection setup, the end station can request the quality of service it needs to suit transmission requirements, and ATM switches will grant the request if sufficient network resources are available. The guaranteed quality of service of cell-based switched access is particularly useful for transporting real-time, interactive communication such as voice or video.

"LAN Emulation allows today's data networking protocol software to enjoy high-speed ATM networking without modification."

The Role of Edge Devices

All ATM traffic-handling decisions are based on destination information in the cell header, not on the content of the cell payloads. In order to move traffic through the ATM network, devices at the boundary or edge of the network convert non-ATM traffic streams into cells. The addition of new traffic types requires only a new edge device, deployed where the demand for such traffic exists.

ATM is a connection-oriented transport service. Within only five bytes of header, an ATM cell cannot carry the full destination address for each cell. Instead it uses an abbreviated address, called a virtual channel identifier, that provides enough information to establish a connection between two ATM stations. Once a connection exists through the ATM network, communications can ensue.

Legacy LANs, on the other hand, employ connectionless transmission technology based on 48-bit addressing. Thus, edge devices must have some way of adapting existing network layer protocols, such as IP and IPX, to the connection-oriented cell-switching paradigm.

Such higher-level details of ATM are being addressed by the ATM Forum, a consortium of vendors, carriers, and users, including 3Com Corporation, formed to expedite industry agreements on ATM interfaces. One very important interface for interoperability with legacy LANs is the LAN Emulation User-to-Network Interface (LUNI). LUNI protocols allow the ATM network and its edge devices to control the virtual connections required for transmission and to emulate the connectionless nature of a LAN. While proprietary LAN Emulation strategies exist, the ATM Forum's LAN Emulation is the only standardized service promoting cross-vendor interoperability.

Abbreviations and Acronyms

AAL
ATM adaptation layer

ATM
Asynchronous Transfer Mode

BISDN
Broadband Integrated Services Digital Network

CCITT
Comité Consultatif International Télégraphique et Téléphonique

DXI
Data Exchange Interface

FDDI
Fiber-Distributed Data Interface

HPSN
High-Performance Scalable Networking

ITU
International Telecommunications Union

LUNI
LAN Emulation User-to-Network Interface

MAC
Media access control

SNMP
Simple Network Management Protocol

SONET
Synchronous Optical Network

UNI
User-to-Network Interface

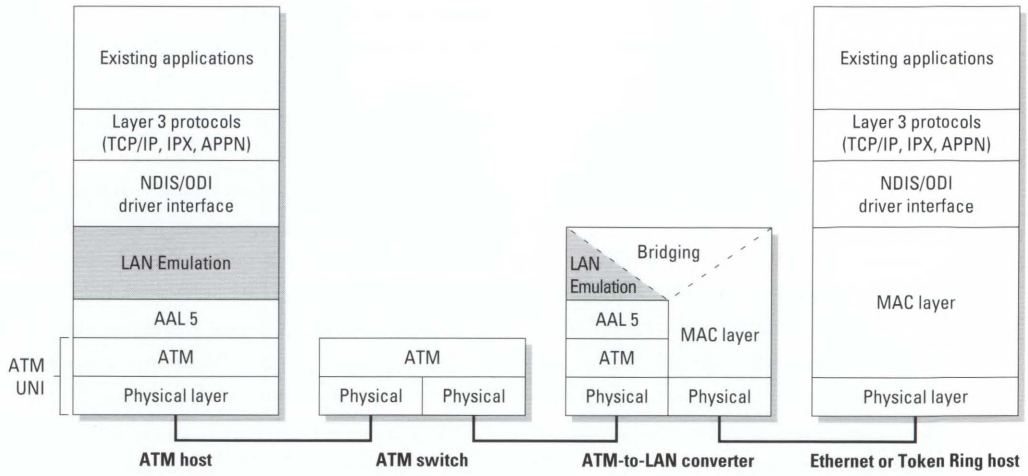

Figure 1. *Conceptual View of LAN Emulation*

Glossary

Asynchronous Transfer Mode
A high-speed, connection-oriented switching and multiplexing technology that can transmit voice, video, and data traffic simultaneously through fixed-length packets called cells.

ATM adaptation layer
A set of protocols that translate user traffic from higher-layer protocols into ATM format.

ATM layer
The part of the BISDN protocol stack that handles most of the ATM routing and processing.

Connectionless communications
A form of packet-switching that relies on global addresses in each packet rather than on predefined virtual circuits.

Connection-oriented communications
A form of packet-switching that requires a predefined circuit from source to destination to be established before data can be transferred.

LAN Emulation
A way for legacy LANs and all higher-layer protocols and applications to integrate transparently with ATM networks.

Switched virtual LAN
A logical network consisting of several different LAN Emulation domains controlled through an intelligent network management application.

Virtual channel identifier
A unique numerical tag for every virtual channel across an ATM interface.

ATM LAN Emulation

According to Version 1 of the ATM Forum LAN Emulation Specification, "The main objective of the LAN Emulation service is to enable existing applications to access an ATM network via protocol stacks like APPN, NetBIOS, IPX, etc., as if they were running over traditional LANs." LAN Emulation is both an edge-device function and an ATM end-system function that allows today's data networking protocol software to enjoy high-speed ATM networking without modification. Legacy end stations can use LAN Emulation to connect to other legacy systems, as well as to ATM-attached servers, routers, hubs, and other networking devices.

A Conceptual View of LAN Emulation

LAN Emulation provides a translation layer between the higher-level protocols of connectionless protocol services and the lower-level, connection-oriented ATM protocols, as shown in Figure 1. Consider the protocol layer differences between the ATM host at the figure's far left and the Ethernet or Token Ring host at the far right. In the BISDN protocol stack, the ATM layer sits directly above the physical layer. Media independence is a driving principle of ATM. Many physical layers are specified, including several for 100 to 155 Mbps. The 155-Mbps WAN interface to the public network carriers will be based on the Synchronous Optical Network (SONET), and other market factors indicate that SONET-based interfaces will predominate across the LAN as well.

The ATM layer manages the header for the ATM fixed-length cell. It accepts the cell payload from a higher layer, appends the header, and passes the resultant 53-byte cell to the physical layer. Conversely, it receives cells from the physical layer, strips off the header, and passes the remaining 48 bytes to the higher-layer protocols. The ATM layer is unaware of the types of traffic it carries, though it does distinguish the quality of service through information learned during connection setup.

The ATM adaptation layer (AAL) sits above the ATM layer. The AAL formats data into the 48-byte ATM cell payload, a process known as segmentation. Once the ATM cells reach their destination, they are reconstructed into higher-level data and transmitted to the respective local devices in a process referred to as reassembly. Because ATM can carry multiple traffic types, several adaptation protocols, each operating simultaneously, exist at the adaptation layer. AAL Type 5 is used for LAN Emulation.

LAN Emulation sits above AAL 5 in the protocol hierarchy. In the ATM-to-LAN converter at the network edge, LAN Emulation solves data networking problems for all protocols—routable and nonroutable—by bridging LAN and ATM addresses at the

media access control (MAC) layer. LAN Emulation is completely independent of upper-layer protocols, services, and applications.

Because LAN Emulation occurs in edge devices and end systems, it is entirely transparent to the ATM network and to Ethernet and Token Ring host devices. LAN Emulation completely masks the connection setup and handshaking functions required by the ATM switch from the higher protocol layers. Conversely, it maps the MAC address–based data networking protocols into ATM virtual connections. The ATM network thus appears to function like a connectionless LAN.

LAN Emulation in Practice

Figure 2 shows how LAN Emulation works in a legacy LAN. Low-end PCs in Ethernet and Token Ring environments access high-end servers with native ATM interfaces through a LAN/ATM switch. Because LAN Emulation makes ATM look like a classical LAN, standard bridging techniques allow the LAN/ATM switch to provide protocol-independent connectivity.

No change is required in the legacy PCs, yet they experience improved performance because of the high input/output capacity of the server, made possible through the high-speed ATM interface. In addition, they benefit from the dedicated bandwidth provided by the switched LAN implementation.

Multiple Emulated LANs

The ATM Forum LAN Emulation standard also supports the implementation of multiple emulated LANs within a single ATM network.

LAN Emulation is implemented through a client-server model. A LAN Emulation client, such as workstation software, resolves MAC addresses into ATM addresses, managed by server functions. Each client connects to the server by a virtual connection. Only those clients connected to the same server can learn about each other and communicate directly. Logically segmenting the network across multiple server functions—which can be stand-alone devices, software in end systems, or ATM switch modules—allows multiple emulated LANs to exist simultaneously on the same physical network.

Figure 3 on page 8 shows the physical and logical view of multiple emulated LANs. In the physical view, the router runs two instances of the LAN Emulation client—M_3 for the marketing group and E_3 for the engineering group. Each departmental server keeps track of its clients through a resident database. When marketing client M_1 sends a packet to engineering client E_1, the marketing server checks its database for a match. Finding none, it maps the MAC address to the router (M_3/E_3), which then forwards the packet to the engineering server for delivery to E_1. The packet travels through the ATM switch twice.

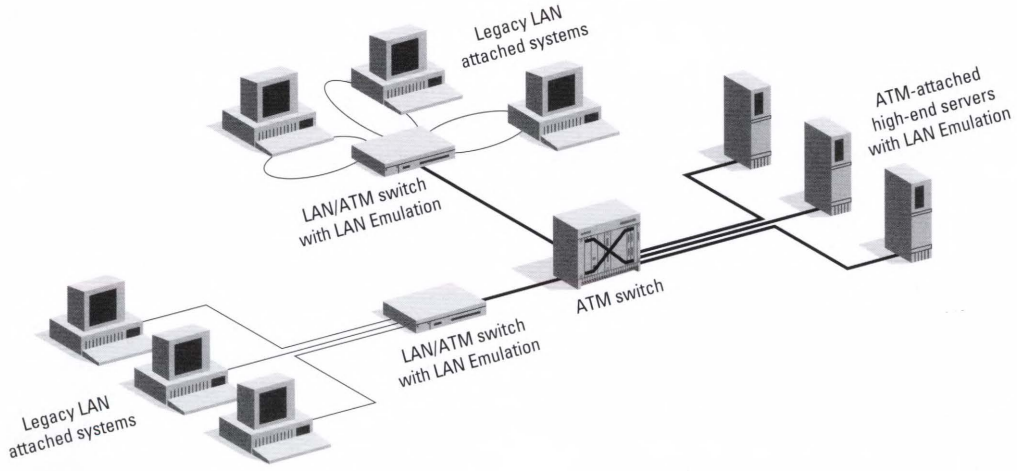

Figure 2. *LAN Emulation for Migrating Legacy LANs*

3Com ATM Products

For information about the features and availability of the following 3Com ATM products, contact your local 3Com sales representative.

- *CELLplex™ 7000 ATM Backbone Switch*
- *CELLplex 7200 Ethernet/ATM Departmental Switch*
- *LinkSwitch™ 2700 ATM Ethernet/ATM Workgroup Switch*
- *NETBuilder II® ATM UNI Module*
- *LANplex® 6000 ATM UNI Module*
- *LANplex 2000 ATM UNI Module*
- *LinkBuilder® MSH™ with ATM LAN Switch Module*
- *ATM Network Adapter*

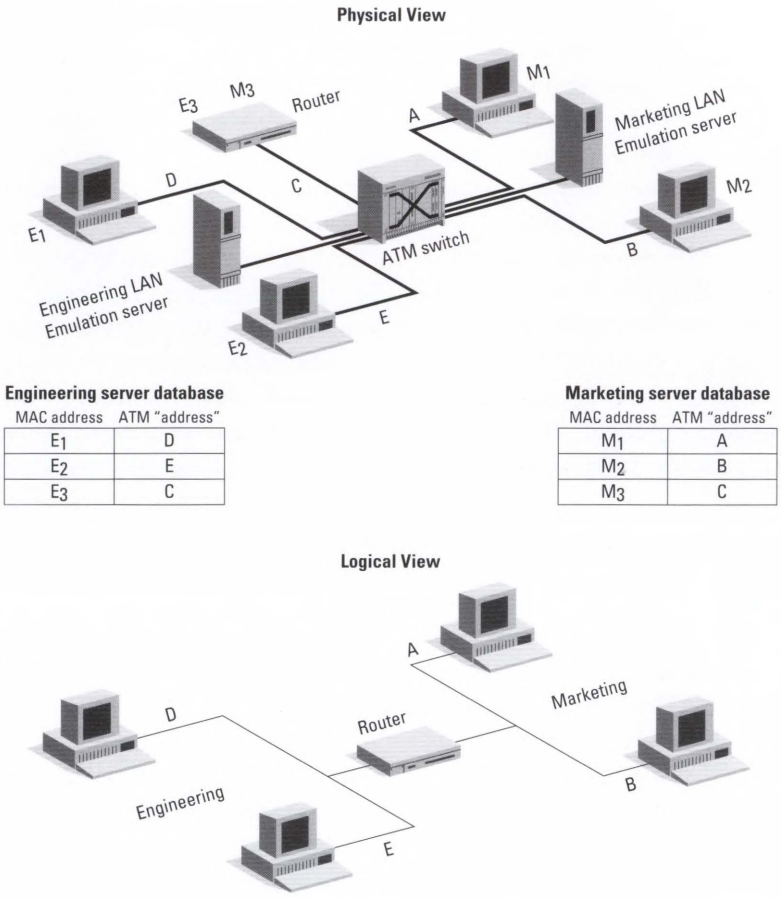

Figure 3. *Physical and Logical Views of Multiple Emulated LANs*

The logical view looks like the physical layout of today's LANs and is consistent with the goal of LAN Emulation. Departmental clients communicate directly with each other, directly with servers, and indirectly with other departments through a network router.

ATM Virtual LANS

The ability to define multiple emulated LANs permits the network manager to create several bridged LAN domains within a single ATM network. (A domain consists of a group of users that can communicate directly.) A switched virtual LAN results when several different LAN Emulation domains exist through one or more switches on a network.

A critical component of managing virtual LANs is a network management application that goes beyond physical connectivity to show how systems logically interconnect across the network. 3Com's Transcend suite of network management applications provides this type of "bigger picture" by allowing administrators to manage their corporate networks as logical groups of related devices.

Virtual LANs create secure workgroups, erect firewalls against broadcast storms, use flow control to make better use of existing network bandwidth, and allow networks to be reconfigured—all without changing cabling or adding equipment. Network administrators can implement adds, moves, and changes simply by redefining groups in the network management system and remotely configuring software in the end device or ATM switch.

The Routing Future of ATM Virtual LANs

LAN Emulation is a coupling of ATM switching with bridging techniques. However, ATM switching must become more tightly coupled with routing to reap the full benefits

of high-speed ATM transmission. The ATM Forum is in the early stages of developing a standard known as Multiprotocol over ATM, which will address issues associated with routing network layer protocols. Under this standard, the routing function as it is known today will be split into two parts:

- The ATM virtual channel will handle the frame-forwarding function.
- The route server function need not reside in any physical location but can be deployed in a number of places to distribute topology and next-hop information.

Consider how this situation affects the scenario in Figure 3. Using LAN Emulation, which works at the MAC address layer, marketing client M_1 must go through the router to communicate with engineering client E_1. When ATM switching is coupled with routing, the route server function can provide enough information to M_1 so that it can connect to E_1 directly, since the devices physically reside on the same ATM switch.

Because all the configuration and software complexity remain in the router, the complexity of configuring the network does not increase. And since forwarding decisions can be based on network layer addresses, more control is available. The result is the speed of ATM with the control of routing.

Scaling Up to ATM

3Com's HPSN architecture can serve as a roadmap for building, growing, and managing network infrastructures. HPSN logically integrates core technologies into the different networking environments in the enterprise network umbrella—building/campus, workgroups, WAN backbones, remote offices, and personal offices.

Most network managers have learned that the safest and most successful changes occur incrementally, where and when they are most needed. The modular HPSN approach to data networking supports this strategy. A natural evolution is to focus on employing ATM with LAN Emulation in the building/campus, workgroup, and WAN backbone environments, letting user requirements dictate the timing and priority.

LAN Emulation also enables high-performance end systems to connect to ATM. The most likely early examples of ATM-attached end systems will be high-end servers. Prices will have to decline substantially before ATM will be widely implemented in personal or remote offices.

An incremental approach to ATM migration combines the best of the various technologies by enhancing existing network investments in three ways:

- ATM deployment extends the life of the installed base of network equipment by boosting performance.
- Combined with such capabilities as LAN Emulation, ATM enhances network management and operations by allowing virtual network configurations.
- Incremental upgrades keep investment and technical risks low while ATM technology matures.

Integrating ATM into Campus and WAN Backbones

In campus and WAN environments, ATM can connect collapsed backbone nodes such as bridges and routers (Figure 4 on page 10). Without changing any vertical risers or distributed end systems, this solution can dramatically increase the aggregate backbone bandwidth. For example, assume that the existing campus backbone network is FDDI. FDDI has an aggregate backbone bandwidth of 100 Mbps, shared by all attached devices.

Integration Checklist for ATM Campus and WAN Backbones

Backbone switch with significant aggregate bandwidth and fail-safe redundancy features, such as hot-swappable interface cards and redundant power supply	✓
ATM interfaces on backbone bridges and routers with ATM Forum LAN Emulation for full multiprotocol support	✓
SNMP-based network management system	✓
Network management application that provides comprehensive view of connectivity across all functional areas of the network	✓

A backbone ATM switch might have 16 ports, each 155 Mbps. Because those 155 Mbps are dedicated to each attached device, the theoretical aggregate bandwidth is 16 x 155, or approximately 2.5 Gbps. The actual aggregate bandwidth will fall somewhere between 155 Mbps and 2.5 Gbps, depending on the traffic pattern.

Figure 4 shows the existing FDDI backbone operating in parallel with the ATM-capable backbone, with both backbones connected to the ATM switch or ATM WAN service. This represents a fail-safe transition strategy. Should the ATM-capable backbone device fail, routers and bridges would automatically switch to the FDDI backbone.

The campus backbone bridge/router in the basement is equipped with an ATM User-to-Network Interface (UNI) for direct connectivity to the ATM switch, while the WAN backbone attaches to the ATM switch through a UNI or ATM Data Exchange Interface (DXI) to a service multiplexer, typically a special CSU/DSU. The service mux allows non-ATM-capable devices to gain access to ATM without a hardware change, and it multiplexes the various traffic streams into the access link to the WAN. When used in conjunction with an ATM DXI, the service mux can perform the AAL segmentation and reassembly functions. This configuration allows wide-area network users to gain access to integrated public network services through a single high-bandwidth link.

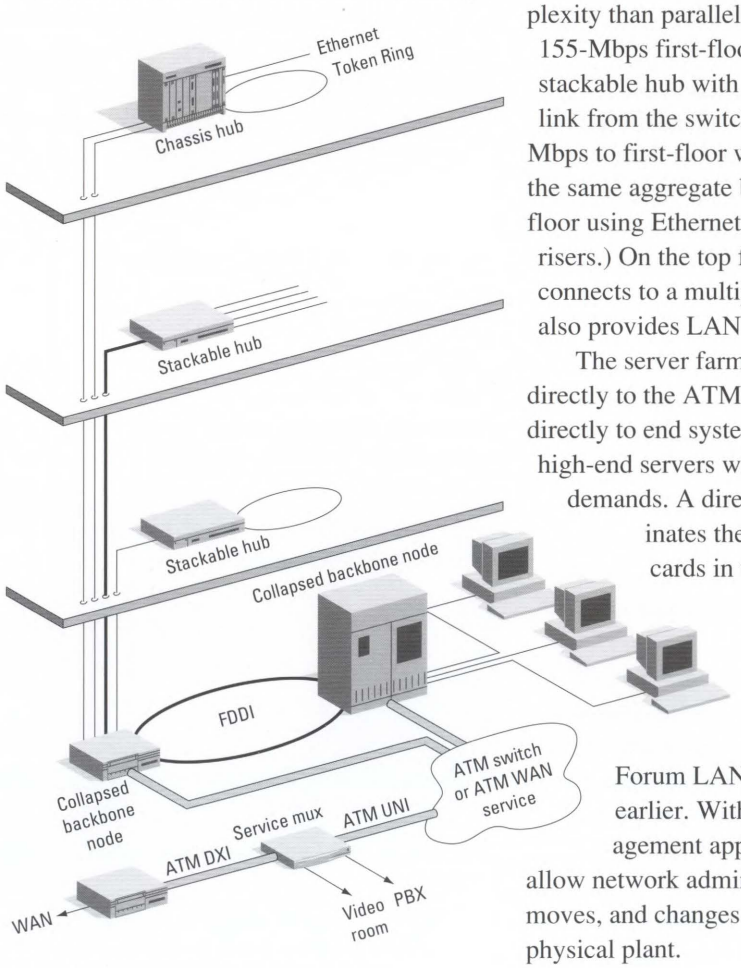

Figure 4. *ATM in the Campus and WAN*

Integrating ATM into the Building Backbone, Workgroups, and Server Farm

Integrating ATM into the building, workgroups, and server farm extends ATM to the wiring closet hubs (Figure 5). High-speed ATM links provide high-bandwidth connectivity to the first and top floors with less complexity than parallel LAN links. The 155-Mbps first-floor link connects to a stackable hub with LAN switching. Each link from the switch provides a dedicated 10 Mbps to first-floor workgroups. (To provide the same aggregate bandwidth to the first floor using Ethernet would require 15 vertical risers.) On the top floor, the ATM link connects to a multipurpose chassis hub that also provides LAN switching.

The server farm in the basement connects directly to the ATM switch. Delivering ATM directly to end systems is especially useful for high-end servers with intensive input/output demands. A direct ATM connection eliminates the need for parallel LAN cards in the server.

Another important benefit of ATM in the building is virtual LAN capability, made possible through ATM Forum LAN Emulation as described earlier. With the right network management application, virtual LANs allow network administrators to make adds, moves, and changes without modifying the physical plant.

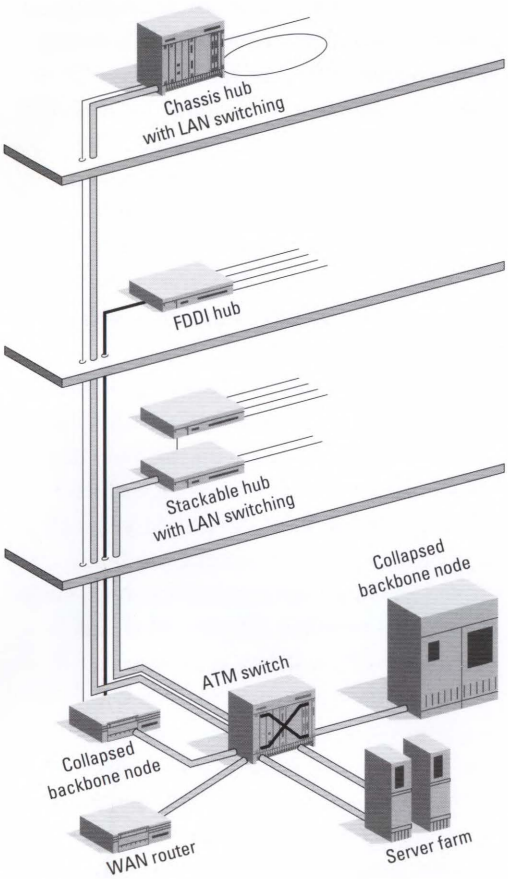

Figure 5. *ATM in the Building Backbone, Workgroups, and Server Farm*

Note that the second floor in this figure remains unaffected by the ATM upgrade. Full connectivity exists between old and new technologies through a bridge/router that complies with the ATM Forum LAN Emulation standard.

Conclusion

ATM is an emerging data networking technology designed for efficient high-speed transfer of any type of traffic. The ATM Forum LAN Emulation standard, approved in early 1995, has paved the way for integrating ATM with legacy systems across the enterprise. Using LAN Emulation, companies can enjoy the bandwidth and quality of service benefits of ATM without modifying existing protocols, software, or hardware. By defining multiple emulated LANs across an ATM network, network managers can create switched virtual LANs for improved security and greater configuration flexibility. Intelligent network management applications like 3Com's Transcend network management suite provide logical views and control of switched virtual LANs.

Phasing ATM into the existing infrastructure by following 3Com's HPSN road map lets companies protect their LAN/WAN equity and benefit from higher performance at reduced risk. By selecting equipment that supports the ATM Forum LAN Emulation standard, network managers can scale up to ATM incrementally, as needs require, for maximum performance with a minimal disruption of ongoing network activity.

Integration Checklist for ATM in the Building, Workgroups, and Server Farm

Backbone switch with significant aggregate bandwidth and fail-safe redundancy features, such as hot-swappable interface cards and redundant power supply	✓
LAN-to-ATM switch with ATM Forum LAN Emulation	✓
ATM adapters with ATM Forum LAN Emulation for high-end servers	✓
ATM interfaces on backbone bridges and routers with ATM Forum LAN Emulation for full multiprotocol support	✓
SNMP-based network management system	✓
Network management application that provides logical/virtual as well as physical views and control of network connectivity	✓

Feature Article

The LANplex 2000 Architecture

Foundation of the New Generation High-Performance Switching Hubs

By Brendon Howe

Brendon Howe is a product manager for the switching division at 3Com Corporation. Prior to joining the company in February 1994, Brendon was senior product manager and tactical marketing manager for Proteon. He has also worked at Digital Equipment Corporation and Analog Devices.

Brendon holds an M.B.A. from New York University and a B.S. in electrical engineering from Boston University.

The LANplex® 2000, the first of 3Com's new family of ASIC-based switching hubs, debuted in January 1995. The hub incorporates a state-of-the-art custom ASIC chip and an innovative dual-processor system architecture. These hardware technologies form the foundation for a host of powerful switching and system management features. The hardware and software design of the hub platform gives it unprecedented price-performance and functionality.

This article describes the Intelligent Switching Engine (ISE) chip and dual-processor system architecture that is the cornerstone of 3Com's LANplex 2000 product line. It describes the implementation of this architecture in the new LANplex 2000 switching hub products, and describes some of their advanced switching features. The article is written for network managers who are interested in learning more about the advanced system architecture of 3Com's newest generation of switching hubs.

The LANplex Intelligent Switching Engine

Application-specific integrated circuit (ASIC) technology has made a significant impact on the networking industry. With this technology, engineers can embed high levels of software functionality, such as switching, bridging, network translation, and some levels of routing, onto specialized microchips. This migration of functionality onto integrated circuits, although time-consuming and capital-intensive to implement, produces tremendous benefits in product cost, performance, and reliability.

3Com actually has three switching ASICs under development. The ISE-chip, discussed in detail in this article, is designed primarily for Ethernet and FDDI switching environments. The ZipChip™, currently under development, is designed for Ethernet and ATM switching environments. The BRASICA (bridging ASIC) chip, also in development, is designed for Ethernet and Fast Ethernet switching environments. These three ASICs will provide customers with LAN technology choices that fit a wide range of network environments.

3Com's ISE-chip is the first graduate of the switching division's ongoing ASIC development. 3Com embarked on ISE-chip development in 1993, taking advantage of its expertise—more than 150 person-years of switching experience—to migrate field-proven hardware and software designs, such as those used in the LANplex 5000 and LANplex 6000 switches, onto silicon.

In the switching hub, the ISE-chip performs preprocessing tasks that accelerate traffic flow, and handles virtually all of the traffic through the device. The sophisticated design of the ISE-chip, which contains over 200,000 gates and implements over 1.2 million transistors, has produced a fast, full-featured switch. The chip will be the foundation for 3Com's Ethernet–FDDI switching products, including the LANplex 2000 and future modules for the LANplex 6000, as well as 3Com's stackable LinkSwitch™ family.

ISE-Chip Switching Performance

ISE-chip technology is capable of forwarding packets at rates in excess of the maximum packet transmission rate when all 18 networks supported by the switch are connected. This maximum transmission rate, which is equal to 565,469 packets per second (pps), is the result of switching nontranslational FDDI and Ethernet streams (all Ethernet traffic switched to Ethernet segments and all FDDI traffic switched to FDDI segments), at the smallest allowable packet size for each network interface. This packet rate is derived from switching eight 64-byte Ethernet streams (through 16 Ethernet ports) while simultaneously switching one 17-byte FDDI stream (through two FDDI ports) through the switch.

ISE-Chip Elastic Packet Buffering

A key to the power of the ISE-chip is its packet buffering capability. Adequate packet buffering is absolutely essential for high-speed performance in switching hubs. Among other things, buffering accommodates the speed mismatch within the switch when it forwards traffic between an FDDI-based server and Ethernet-based clients. Unlike the static buffering technique used by most switches today, the ISE-chip's elastic packet buffering design supports a combination of static and dynamic packet buffering techniques (Figure 1).

Static buffering techniques allocate fixed amounts of buffer memory to every port, varying in most cases only with the speed of the network itself. This buffer space is always guaranteed to each port. The problem is that no port can get more than the allocated amount of buffer space, even for an instant. As a result, the switch may drop packets on ports that are out of buffer space, even though other ports are not using their buffers at all.

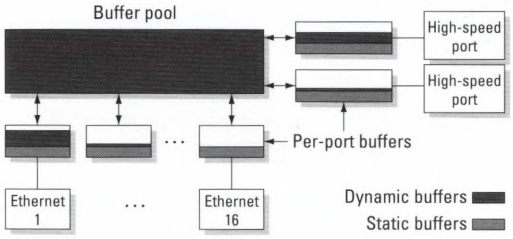

Figure 1. *ISE-Chip Elastic Packet Buffering*

The combination of dynamic and static packet buffering ensures that there is adequate packet buffering for all types of switching demands.

> **"A key to the power of the ISE-chip is its packet buffering capability."**

Dynamic buffering, on the other hand, allocates buffers on an as-needed basis. These buffers are set aside in a shared pool, which can be divided dynamically as congested points come and go in the switch. A flexible allocation of buffer space results in better overall throughput, because it minimizes packet loss. Dynamic buffering allocates more buffer space to the ports that need it, so ports won't drop packets because they are out of buffer memory. However, dynamic buffering often requires more system overhead to manage the buffer allocation and impacts system cost and performance. Another disadvantage is that one exceptionally busy port can consume all available buffer space, leaving none for the other ports.

The ISE-chip's elastic packet buffering technique combines both static buffering and dynamic buffering techniques. By integrating this functionality in the ISE-chip, the LANplex 2000 can support this unique buffering technology with no impact on cost or performance. Like dynamic buffering, it supports large bursts from higher-speed ports to slow ports; like static buffering, it prevents one busy port from starving the other switch ports. Elastic packet buffering is an inherent part of the ISE-chip design. And because the buffering is handled by the ISE-chip and not any additional processors, there is no sacrifice in performance.

LANplex 2000 Hardware Architecture

The ISE-chip is the hardware foundation of 3Com's new LANplex 2000 switching hub product line. Its innovative dual-processor architecture, combined with the ISE-chip, provides the computing engine for the hub's intelligent switching and system management features. Figure 2 on page 14 illustrates the LANplex 2000 system architecture.

The RISC-based packet processor is designed to handle advanced switching functionality: extensive packet filtering, integrated routing, and much more. The processor also maintains comprehensive statistical information—much of which is gathered by the ASIC in real time—such as per-port summaries of packet errors and transmit/receive frame rates.

The dedicated 68340 management processor is designed to handle nearly all of

Abbreviations and Acronyms

ASIC
Application-specific integrated circuit

ATM
Asynchronous Transfer Mode

BRASICA
Bridging application-specific integrated circuit

FDDI
Fiber-Distributed Data Interface

ISE
Intelligent Switching Engine

IEEE
Institute of Electrical and Electronics Engineers

IP
Internet Protocol

IP/RIP
Internet Protocol/Routing Internet Protocol

MIB
Management information base

pps
Packets per second

RISC
Reduced instruction set computing

SNMP
Simple Network Management Protocol

SMT
Station management

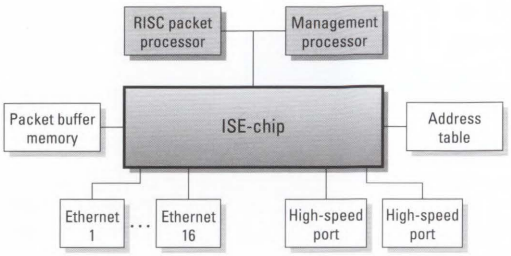

Figure 2. *LANplex 2000 System Architecture*

The dual-processor architecture and the ISE-chip distribute switching and management intelligence and drive switch performance.

the system management functions without degrading switching performance. The dedicated processor can continuously monitor traffic, enable and disable ports, change bridging or routing parameters, download new system software, and restart the switch. The processor also manages bridging parameters such as the Spanning Tree protocol, and routing tables, if routing is used.

The LANplex switch can be managed from either the embedded administrative console or from any SNMP-based management system, and includes support for the following MIBs: Ethernet MIB, FDDI SMT MIB, Bridge MIB, IP Router MIB, MIB II, and the LANplex System MIB. The system also includes an SNMP-to-SMT MIB proxy that allows an SNMP-based management system to access SMT information directly.

Intelligent Switching Software

The dual-processor architecture and ISE-chip in the LANplex 2000 system are the engine that drives performance. To this hardware foundation is added a comprehensive suite of intelligent, integrated switching features:

- Low latency, store-and-forward switching
- User-defined packet filtering
- Virtual networks
- Intra-network IP/RIP routing
- IEEE 802.1d or Express Switching
- IP fragmentation

The sections that follow describe how each of these switching features is implemented in the LANplex 2000 switching hub product line.

Store-and-Forward Switching, Low Latency

The LANplex platform uses store-and-forward packet switching, which performs a complete packet integrity check before forwarding the packet. The integrity check virtually eliminates the chance of propagating an Ethernet "runt" packet or other corrupt data.

The LANplex 2000 system takes advantage of the ISE-chip performance, delivering switching and filtering at full wire speeds with all packet sizes on all ports with extremely low latency. In all cases, switching performance is limited by the packet transmission rate of the network itself, not the packet forwarding rate of the ISE-chip.

"The dual-processor architecture and ISE-chip in the LANplex 2000 system are the engine that drives performance."

User-Defined Packet Filtering

The LANplex system adds a second layer of packet filtering on top of the standard filtering provided by a traditional transparent bridge. This user-defined packet filtering further restricts which packets are forwarded through the bridge. By taking advantage of this powerful feature, network managers can improve network performance by eliminating unnecessary broadcast traffic, provide additional security, and logically segment a network to support virtual workgroups.

The LANplex packet filtering mechanism is very flexible. The user-definable packet filtering language can filter on one or more of the following packet attributes:

- Source or destination MAC address
- Protocol type
- Field type, such as broadcast or multicast

Network managers can define complex filters comprised of many different simple comparisons. Filters can also be applied simultaneously to a single port segment, groups of segments (port grouping), or groups of individual users (MAC address grouping). This flexibility allows network managers to use packet filters for a number of unique applications on the network. Because it is based on a combination of ASIC and RISC design, there is virtually no performance degradation when packet filtering is applied.

Virtual Networks

A virtual network is a group of users that appear to be interconnected, even though they may physically reside on different LAN segments. Virtual networks are often used to connect cross-functional workgroups into a single broadcast domain. This domain can eliminate unnecessary traffic over the backbone, provide security and levels of access, and allow network managers to make changes quickly and easily.

The LANplex switch provides three ways to construct virtual networks: grouping by port, grouping by station addresses, or grouping by routed subnets. (For more information, see "Three Approaches to Implementing Virtual LANs" in the January 1995 issue of *3TECH*.) The LANplex switch provides the flexibility to use all three virtual LAN configurations across the entire enterprise network, allowing network managers to choose the method that best meets their needs.

Intra-Network IP/RIP Routing

The LANplex 2000 system supports IP/RIP intra-network routing within the switch. This capability lets virtual subnets span multiple switched segments within a system or across multiple systems (Figure 3). One of the most common uses of integrated routing is to introduce a switch into an existing IP environment without having to modify the IP subnet assignments. This is accomplished by providing a common IP subnet across multiple switched segments. Traffic within the subnet is switched, while traffic between the different subnets is routed.

IEEE 802.1d or Express Switching Support

The LANplex switch can run in either 802.1d or Express Switching mode to interact with the other switches and routers it is connected to over the backbone.

802.1d, the industry standard for transparent bridging, includes a predefined method for obtaining and storing network addresses. A switch using this bridging mode maintains a local table of every device that is attached to every port on the network. It automatically learns and ages addresses according to the 802.1d standard. The disadvantages of this mode are that it can often produce broadcast traffic when searching for an unknown address, and it requires comprehensive address table management.

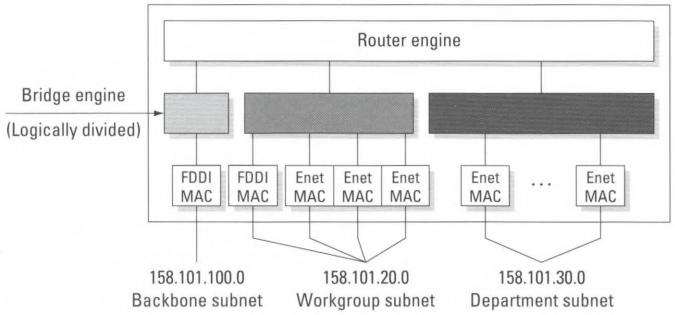

Figure 3. *Intra-Network Routing*

This feature supports multiple ports per routed subnet. IP subnets can span multiple switched segments.

Express Switching, on the other hand, uses a backbone connection to eliminate a high percentage of packet flooding. In Express Switching mode, the address switching table applies to all ports except the backbone port. The switch learns only the local ports; station addresses over the backbone are not learned. As a result, bridging domains (in other words, the number of users) can greatly exceed the address table capacity of the switch. The backbone connection can be configured to be any one of the switched LAN segments.

Both 802.1d and Express Switching bridging modes support the Spanning Tree protocol, which detects network loops and blocks redundant paths to ensure that only one route exists between any two LANs.

IP Fragmentation

IP fragmentation is specified in RFC 791 (Internet Protocol) and RFC 1122. It allows FDDI packets greater than 1518 bytes to be forwarded to stations on an Ethernet segment by segmenting the larger FDDI packets. IP fragmentation also allows FDDI-attached stations to transmit and receive maximum packet sizes (up to 4500 bytes). This feature maximizes network efficiency while maintaining the ability of FDDI stations to bridge to Ethernet-attached stations.

Summary

The LANplex 2000 product line is the first implementation of 3Com's powerful new ISE-chip and dual-processor system architecture. The ISE-chip combines fully integrated switching, packet filtering, and routing acceleration between Ethernet and FDDI networks at full wire speeds. The RISC packet processor provides advanced switching and comprehensive packet monitoring. The management processor handles system management functions, monitors traffic, and manages bridging parameters and routing tables.

These hardware innovations, combined with intelligent switching features, provide the LANplex switch with integrated switching, routing, and bridging functionality. The switching features include store-and-forward switching, user-defined packet filtering, virtual networks, integrated IP/RIP routing, IEEE 802.1d or Express Switching, and IP fragmentation.

The chip will be the foundation for 3Com's Ethernet–FDDI switching products, including the LANplex 2000, future modules for the LANplex 6000, and 3Com's stackable LinkSwitch family.

Feature Article

Using SmartAgent Gauges in LinkBuilder FMS II and LinkBuilder MSH Hubs

By Gordon Hutchison

SmartAgent™ software agents in LinkBuilder® FMS™ II and LinkBuilder MSH™ hubs enable network managers to implement automatic, comprehensive, hub-based network monitoring and self-recovery. These SmartAgent software agents use gauges to monitor traffic levels and error counts to detect exceptional conditions.

This article explains the capabilities and benefits of SmartAgent gauges in 3Com's stackable LinkBuilder FMS II and chassis-based LinkBuilder MSH hubs. It presents four customer scenarios that illustrate how gauges can be used to simplify and automate network management tasks. This article will be of interest to network managers who want to expand their hub management capabilities.

SmartAgent Network Management Software

The SmartAgent software agent is a sophisticated suite of device-based network management software that can support numerous network management and communications protocols. It is a key component of 3Com's Transcend™ network management strategy, which is built on the concept of object-oriented connectivity groups and intelligent management software agents integrated into 3Com's adapter, router, and hub products.

SmartAgent technology locates management intelligence in the software agent running in a managed device. This reduces the computational load on the network management station and helps minimize management-related network traffic. SmartAgent gauges provide the network manager with many advanced capabilities, including the ability to do the following:

- Perform broadcast throttling
- Monitor the ratio of good to bad frames
- Switch a resilient link pair to the standby path if the primary path corrupts frames
- Report if traffic on vital segments drops below minimum usage levels
- Carry out multiple actions at different problem severity levels
- Disable ("blip") a port for five seconds to clear problems, and then automatically reconnect

The SmartAgent's modular architecture separates management agent and protocol interfaces from device-specific network management applications. Because of this device independence, high levels of network management functionality can be incorporated into a wide array of networking products. To date, 3Com has implemented SmartAgent software in EtherLink® and TokenLink® adapters, NETBuilder II® bridge/routers, LinkBuilder FMS II stackable hubs, and LinkBuilder MSH multi-services hubs.

The SmartAgent includes several areas of functionality:

- Management agents and their interfaces to supported network management protocols
- A generic interface to network management protocols
- A generic interface to the device-specific software that operates on MIB items
- A gauges subsystem for system monitoring and detection of exceptional conditions

Figure 1 illustrates how SmartAgent software fits into the overall network management system.

How SmartAgent Gauges Work

Gauges are an integral component of SmartAgent software agents. They provide the

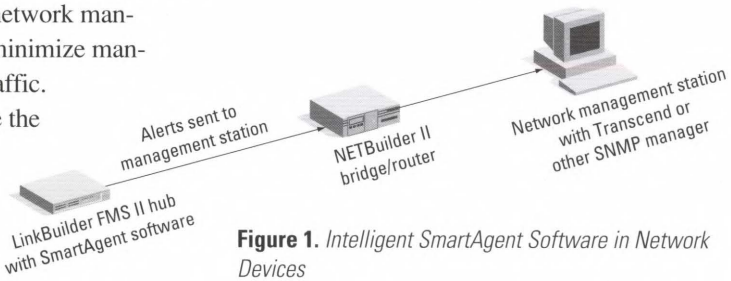

Figure 1. *Intelligent SmartAgent Software in Network Devices*

Gordon is a marketing engineer in the premises distribution division at 3Com Corporation.

Prior to coming to 3Com, Gordon participated in several projects with the New South Wales Public Services. He then joined MM Data Networks, the Australian distributor for BICC Data Networks, later moving to England to join BICC Data Networks as an international technical consultant. When 3Com purchased BDN, Gordon joined the system engineering team.

Gordon holds a B.S. in computer programming from the University of Sydney in Australia.

means for the SmartAgent management agent to continuously monitor routine statistics such as traffic levels or error counts in order to detect exceptional conditions. In fact, gauges provide much of the "smart" in the SmartAgent.

SmartAgent gauges monitor autonomously, without involving the network management station. This is a significant improvement over traditional network management techniques, in which the network management station polls the managed device for information, compares the information it receives with normal values, and issues an alarm or warning when an unexpected value is obtained. In addition to reducing the computational and management load on the management station, gauges also reduce network traffic by removing polling packets that would otherwise be needed to constantly monitor the network.

Network managers can attach a gauge to any counter or integral-type MIB variable. For example, MIB variables such as percentage of bandwidth and rate of corrupted packets provide the network manager with useful information for monitoring the health of the network.

Gauges constantly monitor the value of a counter and generate traps (event messages) when a predefined threshold is exceeded. They can also carry out other actions when the gauge "fires." All responses to exceptional conditions can be configured by the network manager.

Statistics Monitored by Gauges

In addition to reporting counts of simple statistics, SmartAgent software agents help network managers monitor network and device performance by providing meaningful derivatives of raw statistical data.

For example, the raw counters of corrupted packets and alignment errors do not provide information about the number of good frames. So the SmartAgent software agents also report the number of errors per 10,000 frames. This allows a gauge to monitor the rate of errors, rather than just the error number.

Similarly, SmartAgent software agents report on the number of frames, the number of octets, and the bandwidth used. These "smart statistics" enable simple monitoring of relevant information, instead of endless monitoring of raw data. The box entitled "Statistics Monitored by SmartAgent Gauges" shows the full list of statistics that can be monitored by gauges.

Setting Up SmartAgent Gauges to Monitor System Activity

Gauges provide a versatile, general mechanism for setting thresholds and sampling intervals to generate events on any counter or integer maintained by the LinkBuilder FMS II or LinkBuilder MSH hub. These include port traffic statistics, port error statistics, and hub-related statistics. Both rising and falling thresholds may be set, since both can indicate network faults.

For example, crossing a high threshold may indicate network performance problems; conversely, crossing a low threshold may indicate the failure of a scheduled network backup. After a rising threshold is crossed, another rising event is not generated until the matching falling threshold is crossed (Figure 2).

Defining a gauge requires two factors: the sampling time and the number of samples averaged. The sampling time specifies the time between the samples. More frequent sampling allows a quicker response to network faults. The number of samples averaged determines whether the gauges will respond quickly to short fluctuations in the statistic, or respond to general trends observed over a longer period. Averaging multiple samples ensures that the hub does not respond prematurely to samples that are only slightly over the threshold.

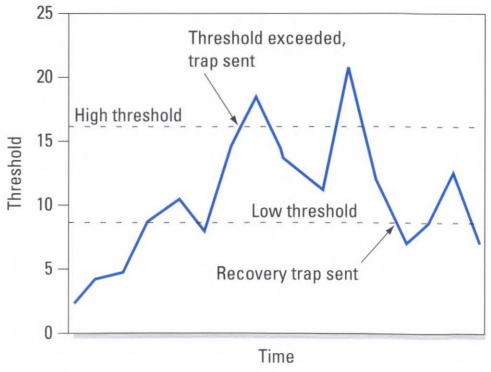

Figure 2. *SmartAgent Gauge Thresholds*

Statistics Monitored by SmartAgent Gauges

Ethernet MIB Object
Port Operational State
Total Partitions
Link State
Total Received Frames
Last Source Address
Received Frames Unicast
Source Address Changes
Received Frames Multicast
Source Address Learning Mode
Received Frames Broadcast
Resilient Link Configuration
Total Received Octets
Standby Port Activated
Received Octets Unicast
Resilient Pair State Change
Received Octets Multicast
Switch Active Port
Received Octets Broadcast

Security Options
FCS Errors
Need to Know
Alignment Errors
Allow Broadcast
Frames Too Long
Allow Multicast
Short Events
Disable Unauthorized Device
Runts
Authorized Station Learn Mode
Transmit Collisions
Authorized Station
Late Events

Token Ring MIB Object
Jabbers (Very Long Events)
Beacon Counts
Data Rate Mismatch
Soft Errors
Auto Partitions
Active Monitor Changes
Total Errors
Topology Table
Received Frames 64 Octets in Length
Active & Standby Monitor Stations
Received Frames 65 to 127 Octets in Length
Station Removal
Received Frames 128 to 255 Octets in Length
Station Line Errors
Received Frames 256 to 511 Octets in Length
Station Internal Errors
Received Frames 512 to 1023 Octets in Length
Station Burst Errors
Received Frames 1024 to 1518 Octets in Length
Station Monitor Errors
Bandwidth Used
Station Lost Frames
Errors per 10,000 Packets
Station Frame Copy Errors
Media Available
Station Receive Congestion
Port Admin State
Station Token Errors
Default Admin State
Station Frequency Errors
Auto Partition State
Station Soft Errors

Figure 3 shows the effect of defining a gauge to use the average of four samples.

The true power and flexibility of the LinkBuilder SmartAgent gauges lies in the large number of actions that can be performed once a threshold is crossed, and in the ability to implement multiple gauges on the same port, monitoring the same statistic. A list of possible actions is shown in Tables 1 and 2 on page 20.

Port control actions are implemented only when monitoring port statistics, and module control actions are implemented only when monitoring module statistics. If the port or module is disabled by the action, all future values of statistics for that port or module will appear to be zero until it has been reenabled. If the action for "exceed high threshold" is to disable, setting the action on "return below low threshold" to enable will always enable the port.

Four Scenarios

This section presents four scenarios that illustrate how SmartAgent gauges in 3Com's LinkBuilder FMS II stackable hubs and LinkBuilder MSH multiservices hubs can be used to cost-effectively monitor and self-correct a network.

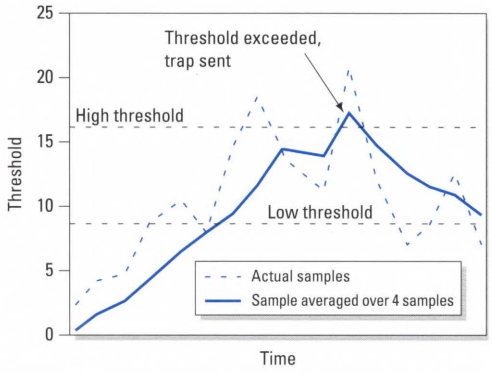

Figure 3. *Defining a Gauge to Average Four Samples*

Table 1. *Exceed High Threshold Actions*

Action Name	Action Description
No action	Not applicable
Notify only	Send trap
Notify and blip port	Send trap, turn port off, wait 5 seconds, turn port on
Notify and blip module	Send trap, turn off all ports on module, wait 5 seconds, return ports to original settings
Notify and disable port	Send trap, turn port off
Notify and disable module	Send trap, turn off all ports on module
Blip port	Turn port off, wait 5 seconds, turn port on
Blip module	Turn off all ports on module, wait 5 seconds, return ports to original settings
Disable port	Turn port off
Disable module	Turn off all ports on module
Notify and switch resilient port	Send trap, then switch to the alternate link

Table 2. *Return Below Low Threshold Actions*

Action Name	Action Description
No action	Not applicable
Notify only	Send trap
Notify and enable port	Send trap and turn port on
Notify and enable module	Send trap, turn ports back to original state
Enable port	Turn port on
Reenable module	Turn ports back to original state

Scenario 1: Monitoring and Evading Broadcast Storms

Company A has a very large, flat network, with a nonroutable protocol used across the entire campus. When broadcast storms occur on the network, the network interface cards in servers become overloaded, resulting in server crashes.

Traditional solution: Implement expensive bridges or switches, which throttle the broadcast storms while continuing to provide for normal levels of broadcast packets. The increased latency results in lower network performance.

SmartAgent gauge solution: Configure the LinkBuilder hubs to monitor every port for broadcast frames. Leave the hub running for a minute, an hour, or a day to establish a baseline for normal broadcast frame traffic. Verify that no problems were reported during the baseline period, and then instruct the LinkBuilder hub to automatically set the gauge to an appropriate level based on the observed baseline (Figure 4). Set the action to "Blip port." When a broadcast storm arises, the affected ports are automatically disabled for five seconds while the storm subsides, and then automatically reenabled.

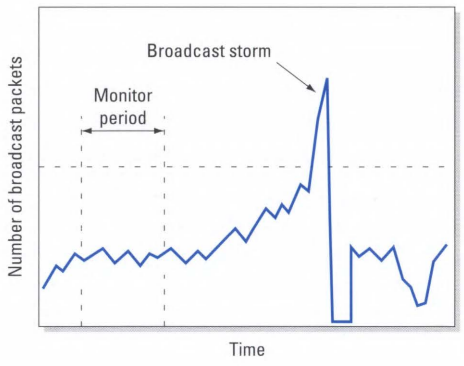

Figure 4. *Using a Gauge to Control Broadcast Storms*

Scenario 2: Incorporating Automatic Resilient Link Switchover

Company B has established an automatic resilient link between two hubs within the same building to provide protection against cable failure. The company wants protection against gradual cable degradation as well as cable breakage.

Traditional solution: Set up an alarm to send a trap to the management station if the number of corrupted packets exceeds a given value. Once the trap has been received, the network manager must manually switch to the backup link, and then use a sniffer to determine whether the ratio of bad to good packets has increased, or whether the total number of packets has simply increased, pushing the total number of bad packets over the limit.

SmartAgent gauge solution: Add a gauge to the primary port of the resilient link pair, which continuously checks the ratio of bad to good packets (Figure 5). If this threshold is exceeded, the hub automatically switches to the standby link and sends a trap to the network management console to log the event. The network manager immediately knows that the primary cable is degrading and that the standby link is being used. The cable can then be automatically switched out of the network until a cable contractor can be called in to fix it.

Scenario 3: Identifying a Faulty Server

An important file server at Company C has been "hanging" mysteriously about once a week. The network manager has had difficulty isolating the problem, since he only learns of the fault well after it occurs. He wants to be notified should the server hang again.

Traditional solution: Set up the management station to constantly ping the server to ensure that it is still alive. The problem with this solution is that the server may not respond to pings, or the ping may be answered by the network interface card even after the server software hangs.

SmartAgent gauge solution: Set up a gauge on the port connected to the faulty server. Configure the gauge to monitor the number of packets sent by the server and received by the port (Figure 6). If the gauge reading falls below the low threshold value for a specified period of time, have the gauge send a trap to the network management station. If the server hangs and no longer transmits packets, an alert is sent to the management station, allowing the network manager to immediately find the fault and take corrective action.

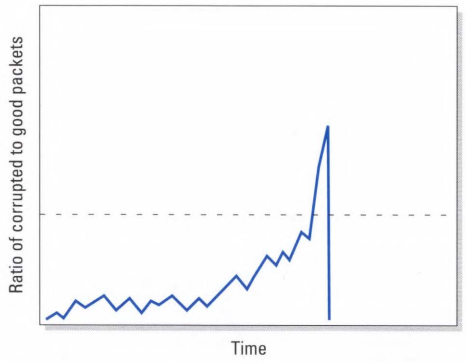

Figure 5. *Gauge Monitoring Ratio of Corrupted Packets to Good Packets*

Scenario 4: Identifying a Fault and Taking Corrective Action

The users in Company D's warehouse have been complaining of slow network response. The problem does not seem to affect users in other buildings. Since all the servers are centralized in the main office block, the network manager suspects a cabling or equipment fault within the warehouse. A network analyzer located with the servers has reported a number of corrupted packets with source addresses of machines in the warehouse.

Traditional solution: Use the network analyzer to progressively trace the source of

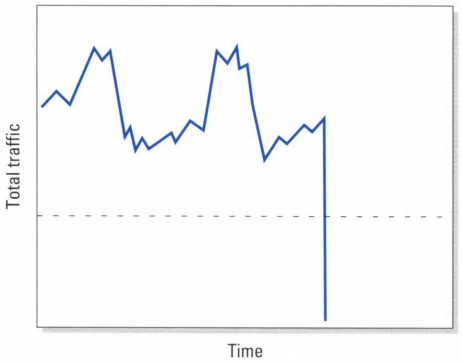

Figure 6. *Gauge Monitoring Traffic Level*

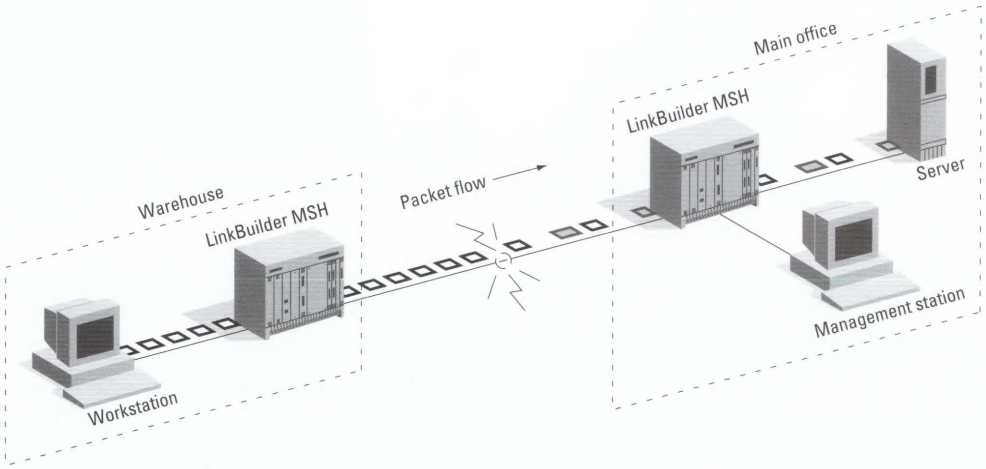

Figure 7. *LinkBuilder MSH Gauge Setup*

the corrupted packets. Once a segment with good packets has been located, the network manager knows that the previous segment or equipment is causing the corrupted packets. This process can require a day or two of testing over the campus network.

SmartAgent gauge solution: Configure all LinkBuilder hubs to send a trap if any error packets are received. The results immediately show that the hub in the warehouse is not seeing any corrupted packets, but the hub in the main office block is seeing a large number of corrupted packets on the port with the link to the warehouse. The network manager has immediately located the problem and can proceed to assess the cost/benefit ratio of replacing the warehouse link (Figure 7).

To quickly establish if the problem is degrading over time, the network manager can use the previously obtained baseline information and reconfigure the gauge to send an alert if the ratio of good to bad packets decreases further.

The network manager can also configure a second gauge, set to a much higher level, to notify and blip the port if the error rate deteriorates. This protects the rest of the network from the packets corrupted by the degrading cable, while not cutting off the warehouse entirely. A third gauge, set at the highest level, can be configured to notify and permanently disable the port should it deteriorate completely.

Summary

3Com's SmartAgent software agents are a key component of the company's Transcend network management strategy. SmartAgent gauges in LinkBuilder FMS II and LinkBuilder MSH hubs allow network managers to monitor network activity and performance. The article has described four scenarios that illustrate how SmartAgent gauges can isolate and monitor system errors to avoid problems, and how gauges can identify and pinpoint system faults and take corrective action. SmartAgent gauges are implemented in the Transcend WorkGroup Manager, Transcend Enterprise Manager, and Transcend Hub Manager network management software for all 3Com LinkBuilder FMS II and LinkBuilder MSH hubs. ◻

Feature Article

Constructing Firewalls

Using the NETBuilder II IP Packet Filtering to Build IP-Based Internet Firewalls

By Bob Konigsberg

Firewalls have a special role on a company network. They allow internal users access to outside services while protecting the network from unauthorized access by outsiders. They also allow the bidirectional flow of data, where authorized, between the public Internet and the private company network.

This article describes how to set up an IP-based Internet firewall on a 3Com NETBuilder II® that allows users on a private network to obtain outside services while protecting private resources. The high-performance NETBuilder II bridge/router can handle the performance challenge that firewall filters present to the system. The article also outlines additional firewall security guidelines.

The simple firewall examples described in this article are not meant to be foolproof mechanisms for protecting your network. The specific firewall configurations that 3Com uses, for example, are more complex than the samples explained here, but are based on the same principles. In addition, the construction of a firewall is no substitute for general network security, in terms of both vigilance and implementing good security procedures on exposed servers and on the network in general.

The Role of Firewalls

As more and more users on private networks want access to outside services such as Telnet, FTP, Internet Mail, and the World Wide Web (WWW), network managers must be able to provide Internet access while protecting their networks from attacks from the outside.

There are three basic classes of attacks against network resources: espionage, or obtaining sensitive information; theft, or the unauthorized use of privately owned resources such as disk space, computing power, and connections; and denial of service, or router or host/server manipulation that damages file systems or restricts legitimate users from network services. Firewalls protect a private company network connected to the Internet from these outside attacks.

Firewalls are typically constructed on a host machine and a multiprotocol router. The host machine, usually UNIX, has two network connections, one on the internal side and the other on the Internet side. This host is configured to prevent it from being compromised by outsiders and to provide detailed logging of system activity for security monitoring. This UNIX host frequently serves as a public server for FTP, e-mail, or the WWW.

The multiprotocol router has extensive filtering capabilities to limit the type and direction of traffic that passes through it. The router is generally not the subject of attack itself other than as a barrier to other, more desirable targets, or as the basis of a denial-of-service attack. This article discusses the router-based portion of firewall security (Figure 1).

Constructing an IP-Based Internet Firewall

This section describes how to construct an IP-based firewall on a NETBuilder II bridge/router. The procedure involves the following seven steps:

- Set up blanket filtering to deny everything not specifically permitted.
- Determine which services to allow through the firewall.

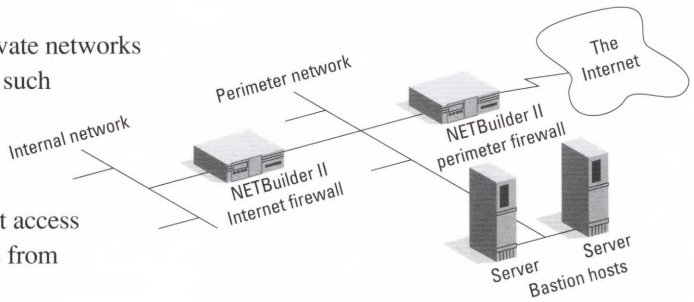

Figure 1. *Basic Network-to-Internet Firewall Design*

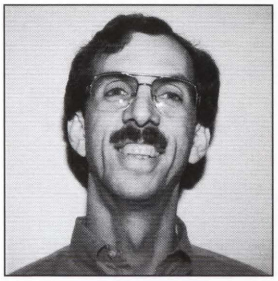

Bob Konigsberg is a senior network engineer in 3Com's MIS division, where he administers the WAN, including the Internet. Before coming to 3Com eight years ago, he tested operating systems at the International Bureau of Software Tests.

Bob has a B.S. in industrial technology from California Polytechnic State University at San Luis Obispo.

Glossary

Firewall
A router or workstation with multiple network interfaces that controls and limits specific protocols, types of traffic within each protocol, types of services, and direction of the flow of information.

Perimeter network
A small, single-segment network between a firewall and the Internet for services that the organization wants to make publicly accessible to the Internet without exposing the network as a whole.

Bastion host
A machine placed on the perimeter network to provide publicly available services. Although secured against attack, it is assumed to be compromised because it is exposed to the Internet.

Secure logging
A method whereby an audit trail of system activity is received from a bastion host and placed in a secure location.

Spoofing
The use of a forged IP source address to circumvent a firewall. The packet appears to have come from inside the protected network, and to be eligible for forwarding into the network.

Firewall Security Guidelines

When you set up a firewall, you should secure the entire perimeter network and its routers (firewalls and others). The following guidelines, although a good starting point, by no means constitute a complete list of security precautions. Other precautions are dictated by the specific nature of the network, as well as by the type and number of servers exposed to the Internet and the types of services allowed through the firewall.

- Set up passwords, preferably long ones with legal characters other than A–Z, for the firewall routers.
- Set up a user password that prevents anyone from observing filter configurations (NETBuilder version 7.0 and later).
- Make sure firewall router passwords are different from passwords on all other routers on the private network.
- Make sure firewall routers learn routes only from certain other trusted routers.
- Make sure the entire perimeter network, bastion hosts, firewall routers, and hubs are physically secured and isolated from the rest of the private network.
- Make sure SNMP commands are accepted only from trusted (internal) hosts, and traps are generated for SNMP commands issued from nontrusted hosts.
- Do not configure PUBLIC, PRIVATE, and ANYCOM community strings at any time. Alternately, SNMP can be turned off if not needed; however, this is not recommended if you have SNMP monitoring capability.
- Deny both Telnet and remote management access on all firewall routers. If you want Telnet access to the routers, connect a communications (terminal) server to the console port. This has the added advantage of autodisconnection during inactivity.

- Select the type and direction of traffic through the firewall.
- Set up filters and filter addresses to allow selected traffic through the firewall.
- Set up a filter to prevent IP source address "spoofing."
- Set up a filter to allow access to the bastion host without compromising network security.
- When all filters are set up, test the firewall in both the permitted and prohibited directions to make sure that it operates as intended from either side.

Step 1: Set Up a Blanket Filter

The two basic approaches to firewall security are: "Everything not specifically permitted is denied," and "Everything not specifically denied is permitted." Depending on the application, router-based firewalls can use both schemes.

For IP-based firewalls, the first approach is the safest. Therefore, the first step is to set filtering to discard all packets, regardless of type, by setting the IP Control to Filtering with the FilterDefAction (Filter Default Action) set to Discard:

```
setd -ip CONTrol=Filtering
setd -ip FilterDefAction=Discard
```

Step 2: Decide Which Services to Allow Through

Decide what services to allow through the firewall, and to which parties. Telnet and FTP services were chosen for the sample firewall in this article, although the same principles apply to Simple Mail Transport Protocol (SMTP – Mail), Domain Name Service (DNS – Name Service), Hypertext Transport Protocol (HTTP – World Wide Web), Gopher, Archie, and others.

Step 3: Define Ingoing and Outgoing Firewall Traffic

Select the type and direction of traffic through the firewall. The general objective of an Internet firewall is to allow users inside a private network to have outbound access, while restricting outside users from inbound access. The types of incoming and outgoing traffic can be identified as follows:

- Inside-originated request to outside service
- Outside reply to the inside-originated request
- Outside-originated request to inside resource
- Inside reply to the outside request

The firewall is constructed to accept the first two types of traffic and deny (or discard) the last two types. Filters are set up to allow inside users to initiate a connection to the outside world and to receive a reply from the outside server. A machine on the outside, however, cannot initiate a call into the private network, nor receive a reply.

Table 1 shows the resulting set of port number and IP address combinations using Telnet and FTP sessions. The source port of the outgoing packet has a user range port number of 1023 or greater. The destination port of the outgoing packet has a Telnet port number (23), an FTP control port number (21), or an FTP data port number (20). The first pair of packets in each group must be able to cross the firewall to the Internet and return from the Internet to the private network.

Step 4: Set Up Firewall Traffic Filters

The first filter set up in this sample firewall discards all traffic packets. Two filters and two FilterAddr (filter address) definitions are now added to allow selected traffic through the firewall (Table 2). All values supplied are in hexadecimal: 23 decimal is 17 hexadecimal, 1023 decimal is 03ff hexadecimal. (If you are not familiar with IP filtering on a NETBuilder II, refer to the "IP Packet Filtering on a NETBuilder II" box on page 26.)

Table 1. *Sample Port Number and IP Addresses for Telnet and FTP*

Source Port	Destination Port	Source IP Address	Destination IP Address	Flow of Information
>1023	=23	Inside	Outside	Telnet request to outside
=23	>1023	Outside	Inside	Telnet reply to inside
>1023	=23	Outside	Inside	Telnet request to inside
=23	>1023	Inside	Outside	Telnet reply to outside
>1023	=21	Inside	Outside	FTP-Ctrl request to outside
=21	>1023	Outside	Inside	FTP-Ctrl reply to inside
>1023	=21	Outside	Inside	FTP-Ctrl request to inside
=21	>1023	Inside	Outside	FTP-Ctrl reply to outside
>1023	=20	Inside	Outside	FTP-Data request to outside
=20	>1023	Outside	Inside	FTP-Data reply to inside
>1023	=20	Outside	Inside	FTP-Data to inside
=20	>1023	Inside	Outside	FTP-Data reply to outside

Table 2. *Filters to All Selected Traffic Through Firewall*

```
add !1 -ip fi %00: = %0017, %02: > %03ff    Tests for Telnet server reply to client
add !2 -ip fi %00: > %03ff, %02: = %0017    Tests for client Telnet request to server
add !3 -ip fi %00: = %0014, %02: > %03ff    Tests for FTP-Control server to client
add !4 -ip fi %00: > %03ff, %02: = %0014    Tests for client to FTP-Control server
add !5 -ip fi %00: = %0015, %02: > %03ff    Tests for FTP-Data server to client
add !6 -ip fi %00: > %03ff, %02: = %0015    Tests for client to FTP-Data server

add -ip fa 127.0.0.0 / 0.255.255.255 < ALL forward tcp 1
add -ip fa 127.0.0.0 / 0.255.255.255 < ALL forward tcp 3
add -ip fa 127.0.0.0 / 0.255.255.255 < ALL forward tcp 5
add -ip fa 127.0.0.0 / 0.255.255.255 > ALL forward tcp 2
add -ip fa 127.0.0.0 / 0.255.255.255 > ALL forward tcp 4
add -ip fa 127.0.0.0 / 0.255.255.255 > ALL forward tcp 6
```

Abbreviations and Acronyms

DNS
Domain Name Service

FTP
File Transfer Protocol

HTTP
Hypertext Transfer Protocol

ICMP
Internet Control Message Protocol

IP
Internet Protocol

SMTP
Simple Mail Transfer Protocol

SNMP
Simple Network Management Protocol

TCP
Transmission Control Protocol

UDP
User Datagram Protocol

WWW
World Wide Web

IP Packet Filtering on a NETBuilder II

IP filtering is done in one or two stages on a NETBuilder II, depending on how precise the filter needs to be. The first stage creates IP FilterAddr (filter address) definitions that filter IP traffic at the address, direction, and protocol levels. For example, with IP filtering enabled and the default filter action set to discard all packets, one conversation could be allowed through the firewall with the following filter:

```
add -ip fa 127.0.45.129 <> 127.0.34.76 forward
```

This filter allows the bastion host on the sample network to conduct any IP traffic with one device inside the network. This is reasonably secure, but it severely limits which internal machines can use the external server.

Wildcard IP Address Fields

Access can be expanded by making one of the IP address fields "wild" with respect to a partial IP address. The following FilterAddr is wildcarded:

```
add -ip fa 127.0.0.0 / 0.255.255.255 <> 127.0.34.76 forward
```

Note that the pattern of 127.0.0.0 is paired with a mask of 0.255.255.255. This means that the address fields in the mask that are set to 0 must exactly match the packet being examined in order to have that FilterAddr applied to it. Fields that are set to 255 may have any value within them. The bit mask used here shows which bits must be matched exactly (set to 0) and which bits may be any value (set to 1). This field may be set anywhere along the 32-bit length of an IP address for precise control over which part of the address is allowed.

This is not very secure if the bastion host becomes compromised, but it would work. To make this more secure, you can modify the filter slightly to eliminate all IP-based protocols except TCP:

```
add -ip fa 127.0.0.0 / 0.255.255.255 > 127.0.34.76 forward TCP
```

At this point, the bastion host (127.0.34.76) can no longer send UDP, ICMP, or any other packets that are not TCP. But there's still more to do to make this secure.

Adding Filter IDs

The second stage in NETBuilder IP filtering specifies particular values within each packet or section of packet. If a filter (not a FilterAddr) is defined that specifies that the TCP source port must be greater than 1023 (most client-based applications use this range), *and* that the TCP destination port must be equal to 23, and this requirement is added to the filter addresses above, the filter then limits the firewall to allow only TCP Telnet traffic from the inside network in general to the bastion host in particular. Since the source and destination ports switch for traffic to and from the bastion host, two filters are specified, one for each direction:

Step 1. Create the filter ID definitions:

```
add !1 -IP Filters %00:>%03ff, %02:=%0017
add !2 -IP Filters %00:=%0017, %02:>%03ff
```

Step 2. Apply the filter definition to the filter address:

```
add -ip fa 127.0.0.0 / 0.255.255.255 > 127.0.34.76 forward TCP 1
add -ip fa 127.0.0.0 / 0.255.255.255 < 127.0.34.76 forward TCP 2
```

This is not as hard as it might look. Most of the offset filtering on a firewall is concerned with the source and destination ports, which are offsets 0 and 2 respectively for both TCP and UDP packets. The filter address will apply the filter definition to whichever portion of the packet the protocol is specified for. For example, to filter TCP/IP traffic, the protocol is specified as TCP, as shown above. To filter (forward in this case) UDP traffic, UDP would be used in the FilterAddr definition.

Looking at the filters and FilterAddr definitions above, notice that, for safety, a computer inside the firewall is allowed to set up a Telnet session with the bastion host, but the bastion host is *not* allowed to do the same thing. It can only reply to inside inquiries, not initiate a session on its own. This pair of filters and FilterAddrs allows only one-way session setup, but two-way traffic can be initiated by inside hosts.

The FilterAddr (-ip fa) definitions in the second set of statements instruct the NETBuilder II to forward only those packets that are both:

- TCP (Not UDP, ICMP, or any other IP-based protocol)
- From a private network client to an outside server or the corresponding reply (127.0.0.0 is the private network address)

Any packets not matching this specific filter set are discarded. There are now enough holes in the firewall to allow Telnet and FTP traffic from the inside out, but not from the outside in. For example, a ping will not go through this firewall in either direction, since there are no specific filters to let Internet Control Message Protocol (ICMP) messages through. ICMP Echo (ping) can be used by an intruder to learn more about the structure of your network. Ping can also be filtered, but it is generally safer to deny anything that doesn't meet a specific need.

Since the FilterAddr definitions specify TCP, the byte offsets shown in the filters themselves are applied to the TCP section of the IP packet. This makes filtering simpler than it might first appear because you don't have to perform extensive calculations about where in the packet a given value might appear. For TCP and UDP, which cover most needs, offset 0 is the source port and offset 2 is the destination port. The decimal and hexadecimal values for some common services are shown in Table 3. Others are available from the /etc/services file on most UNIX hosts and commercial packages for TCP/IP.

When you specify filters, be sure to use the full field width (two bytes in this case) for the port numbers, since the field width in the filter specifies how many bytes will be actually compared for filtering. If only one byte is used, any packet matching that one byte would be permitted through, and more traffic than intended is allowed through the filter.

Step 5: Set Up a Filter to Counteract Spoofing

In the previous steps, filters were used to restrict normal usage. Step 5 protects the network from specific attacks. Spoofing is the use of "forged" source addresses that make an outside packet appear to have come from the inside network so that the firewall will allow it to have access to the private network. Spoofing works on the principle that routers perform route lookups only on the destination IP address in each packet, and pay no attention to the incoming interface of the source IP address; they do not check to see where the packet originated.

Table 3. *Commonly Used IP Service Ports and Protocols*

Service	Decimal Port No.	Hexadecimal Port No.	Protocol Name
FTP Data	20	14	TCP
FTP Control	21	15	TCP
Telnet	23	17	TCP
SMTP (Mail)	25	19	TCP
DNS (Name)	53	35	TCP & UDP
Gopher	70	46	TCP
WWW/HTTP	80	50	TCP

To prevent spoofing, set up a FilterAddr to examine the source address of any packet and discard any packets that appear to come from inside the private network, as shown below:

```
add -ip fa 127.0.0.0 / 0.255.255.255 <>
      127.0.0.0 / 0.255.255.255 discard
```

After all, what would a private machine be doing out on the Internet? Since any packets with an inside source and destination address have no business out on the perimeter network or on the Internet in the first place, they should be discarded.

Step 6: Set Up a Filter to Protect the Network from the Bastion Host

A bastion host is an Internet server placed on the perimeter network to provide services to the public network. By virtue of its exposure, it is considered to be compromised and unsecured, and in need of protection. The ideal protection is to have a Class C address separate from the primary network. The next best protection is to obtain a tiny subnet from your

Additional Reading

For references on firewall construction, pitfalls found when building firewalls, and types of attacks to expect from hackers, you may want to read the following:

Firewalls and Internet Security: Repelling the Wily Hackers *by Cheswick and Bellovin. Menlo Park, CA: Addison-Wesley, 1994.*

"Network (In)Security through IP Packet Filtering" by Brent Chapman, FTP@GreatCircle.COM /pub/pkt_filtering.ps.Z

Internet supplier. However, if you don't have either advantage, you can still protect the bastion host by using a small subnet from your primary network and adding specific filters to allow inside personnel to maintain the bastion host without compromising network security.

For example, assume that the bastion host's address on this sample network is 127.0.34.76. Using the same filters as defined for the firewall itself (including anti-spoofing filters), the following FilterAddr definitions are added to allow hosts inside the firewall access to the bastion host, while denying the bastion host access to those inside hosts:

```
add -ip fa 127.0.0.0 / 0.255.255.255 <
      127.0.34.76 forward tcp 1
add -ip fa 127.0.0.0 / 0.255.255.255 <
      127.0.34.76 forward tcp 3
add -ip fa 127.0.0.0 / 0.255.255.255 <
      127.0.34.76 forward tcp 5
add -ip fa 127.0.0.0 / 0.255.255.255 >
      127.0.34.76 forward tcp 2
add -ip fa 127.0.0.0 / 0.255.255.255 >
      127.0.34.76 forward tcp 4
add -ip fa 127.0.0.0 / 0.255.255.255 >
      127.0.34.76 forward tcp 6
```

Bastion hosts should be considered an integral part of the firewall, and be accorded some degree of trust—for e-mail for example—because you have time and money invested in their configuration. Additional protection for bastion hosts can be implemented by setting the Perimeter Net Router to filter all packets, but with the FilterDefAction set to Forward. Specific filters would then be added to discard traffic that you do not want sent to your bastion hosts. Typical examples are:

- No Telnet access to anything from the Internet
- No SMTP mail to a Web (HTTP) or FTP server
- No FTP or Web (HTTP) traffic to a mail server

Using the same filters defined in previous Internet firewall examples, two FilterAddr definitions are added to deny Telnet traffic through the perimeter firewall, as follows:

```
add -ip fa 127.0.34.0 /
      0.0.0.255 ALL discard tcp 1
add -ip fa 127.0.34.0 /
      0.0.0.255 ALL discard tcp 2
```

FTP, SMTP, and HTTP examples would look similar to these instructions. Note that securing the bastion host itself is outside the scope of this article, but can be done by a competent UNIX system administrator.

Step 7: Test the Firewall

When all filters are in place, test the firewall in both the permitted and denied directions to make sure that it operates as intended from both sides.

Summary

Firewalls protect a private company network that is connected to the Internet from outside attacks against network resources.

To construct a firewall, a filter is first set up to discard all packets. After the flow of protocol traffic is carefully analyzed, additional filters are added to allow specified traffic through the router. Anti-spoofing filters are also set up for all advertised network numbers, and preferably all network numbers. Filters are also set up to protect inside hosts from the unsecure bastion host on the perimeter network. When all filters are in place, the firewall is tested to make sure that it operates as intended from either side. If the network has more than one IP address, all network numbers (allowed outside) must be accounted for on the firewall.

If you are careful in your selection of filters, the NETBuilder II provides room for expansion. NETBuilder 7.x and 8.x allow for 64 filters and unlimited filter address sets. ☐

This past January, the Computer Emergency Response Team (CERT) issued a security warning describing IP spoofing. In the next issue of *3TECH*, an article on Internet security will describe 3Com's response to the CERT advisory and will describe in more detail how to configure a NETBuilder II to counteract this type of attack.

Tech Tips

Tech Tips has been developed to help you make your networks more efficient. We try to find the most useful technical tips from a variety of sources including 3ComFacts℠, 3Com's interactive fax service, and Ask3Com℠, 3Com's bulletin board services available on CompuServe®. Both feature technical articles, product and service information, patches, fixes, and utilities. Let me know how you like these tips—are they meeting your needs? I can be contacted via fax at 408-764-5001 or via the Internet at: Suzanne_Dowling@3mail.3com.com

Special thanks to Sergio Arzate, Maria Carattini, Gorden Jennings, Kathy Laymon, Romer Maiel, and Brad Turner for their contributions to Tech Tips.

Installing the TokenLink III 16/4 PCMCIA Adapter with IBM's LAN Support Program Drivers

The following are suggestions to help you install 3Com's TokenLink® III 16/4 PCMCIA adapter with IBM's LAN Support Program (LSP) drivers.

Requirements Prior to Installation

Before you begin the installation, be sure that you have the following:

- TokenDisk® diskette v1.1 with the LANSUP subdirectory. The LANSUP subdirectory contains a connectivity enabler (LSPEN.3CM) and a protocol manager driver (PROTMAN.DOS). Connectivity enablers are small programs that interface directly to the PCMCIA controller and initialize the adapter. Since the LSP drivers do not provide this capability directly, you must use connectivity enablers. The connectivity enabler and protocol manager driver must be loaded before the LSP drivers in your CONFIG.SYS file. You must also create a PROTOCOL.INI file to define the adapter's parameters. If the LANSUP subdirectory is not available on your TokenDisk diskette, retrieve the 3C689N.EXE file from one of the following systems:
  - 3Com BBS (408) 980-8204
  - Ask3Com Forum (CompuServe)
  - 3Com FTP Server ftp.3com.com @ IP address 129.213.128.5
- 16 KB of free upper memory
- IBM LAN Support Program, version 1.31 or higher (version 1.33 or higher is recommended)
- 1 to 1.5 MB of hard disk space

Installation Instructions

Follow these steps to install the adapter:

1. Install the TokenLink III 16/4 PCMCIA adapter in any available slot and attach the cable to your network.
2. Insert the TokenDisk diskette into your diskette drive and start the Configuration and Diagnostics Program. At the DOS prompt, type:

   ```
   A:3C689CFG <ENTER>
   ```
3. Select the Configure Adapter menu item.

   a. Write down the starting Shared RAM address, for example, D0000H. The LSP drivers use a default shared RAM address of D8000H. If the adapter is configured to another shared RAM address, a parameter must be set when loading the driver.

   b. The LSP drivers only support Interrupt Request levels 2 or 9. If the IRQ is not 2 or 9, scroll to this parameter and press <ENTER>. Select IRQ 9 and press <ENTER>.

   Note: If Card Services is installed in your system, you cannot configure this parameter. The connectivity enabler restricts the IRQ to 9.

c. If the ring speed is incorrect, scroll to this parameter and press <ENTER>. Select the appropriate ring speed and press <ENTER>.
d. Tab to <OK> to save the new settings.
e. Exit the Configuration and Diagnostics Program.
Note: During the configuration process, a PROTOCOL.PCM file and a NET.PCM file are created and written to the root directory of the fixed disk drive. These files contain the adapter's configuration settings.

4. Insert the LAN Support Program disk in the floppy drive. Invoke IBM's LAN Support Program Installation Aid. Type:
```
DXMAID <ENTER>
```
5. When prompted for the TARGET FOR LSP, enter C:\LSP3COM. The DXMAID program will copy the necessary LAN Support Program drivers to the LSP3COM directory.
6. Bypass the section that asks to process a driver diskette at the A:\DOS prompt by pressing <ESC>.
7. A new menu appears asking you to select an adapter driver and a protocol. Choose IBM TOKEN-RING ADAPTERS (DXMC0MOD.SYS) and press F4 to save the selection. Press F4 again to continue the installation. The DXMAID program modifies your CONFIG.SYS file to add the LSP drivers.
8. After the installation is completed, edit the CONFIG.SYS file to include device driver statements for the protocol manager driver and connectivity enabler. Also, add the Shared RAM address parameter to the DXMC0MOD.SYS device driver statement. For example, if the shared RAM address is D0000H, your CONFIG.SYS file should contain the following statements:
```
DEVICE=C:\LSP3COM\PROTMAN.DOS/I:C:\LSP3COM
DEVICE=C:\LSP3COM\LSPEN.3CM
DEVICE=C:\LSP3COM\DXMA0MOD.SYS 001
DEVICE=C:\LSP3COM\DXMC0MOD.SYS N ,D000,0,0,0
```
Note: The device driver statements must appear in the order shown.
9. Copy the PROTMAN.DOS and LSPEN.3CM files from the LANSUP directory of the TokenDisk diskette to the C:\LSP3COM directory. Type:
```
COPY A:\LANSUP\*.* C:\LSP3COM <ENTER>
```
10. Copy the PROTOCOL.PCM file to the LSP3COM directory. Type:
```
COPY PROTOCOL.PCM C:\LSP3COM <ENTER>
```
11. Edit the file to include sections for the protocol manager and the connectivity enabler. Save the file as C:\LSP3COM\PROTOCOL.INI.
If Card Services is not installed on your system, the PROTOCOL.INI file should look similar to the following:
```
[Protocol Manager]
   Drivername = Protman$
[LSPEN]
   Drivername=TLPC3$
   PRIMARY
   INTERRUPT=9
   MMIO=0xCC00
   RAM=0xD000
   RAMSIZE=16
   RINGSPEED=16
```

Training Information

To order a course catalog or for information on scheduling, course content, or prerequisites, call 800-876-3Com, press option 7, select 1. Outside the United States and Canada, refer to the 3Com Sales Offices Worldwide listing (inside back cover).

If Card Services is installed in your system, the PROTOCOL.INI file should look similar to the following:

```
[Protocol Manager]
   Drivername = Protman$
[LSPEN]
   Drivername=TLPC3$
   RINGSPEED=16
```

Installing the TokenLink III 16/4 16-Bit ISA Adapter in a Gateway P5-60 Personal Computer

To ensure successful operation of the 3Com TokenLink III 16/4 16-Bit ISA adapter in a Gateway P5-60 computer, you must exclude the adapter's BIOS MMIO memory area and Shared RAM areas from the system's shadow memory area. If you do not exclude enough shadow memory, an MMIO or Shared RAM error will occur while running the adapter's diagnostics. To exclude the adapter's memory areas, you must change the computer's Advanced CMOS Setup.

The default memory configurations for the TokenLink III 16/4 16-Bit ISA adapter are as follows:

	Address	RAM Size
BIOS MMIO	CC000–CDFFF	8 KB
Shared RAM	D8000–DBFFF	16 KB

Using these defaults, change the Gateway P5-60's Advanced CMOS Setup to the following:

DISABLE SHADOW MEMORY SIZE 64K

DISABLE SHADOW MEMORY BASE CC000H

ISA IRQ 9 USED

DISABLE SHADOW MEMORY SIZE is the total amount of memory to be excluded. For the example above, this includes the BIOS MMIO memory area, the Shared RAM area, and all the memory between the two areas. If you configure the MMIO area and the Shared RAM area for contiguous addresses, there will be less memory to exclude and more shadow memory available. A list of available addresses can be found in Appendix B of the *TokenLink III 16/4 16-Bit ISA Adapter User Guide*. To change the default settings, use the TokenLink III Configuration and Diagnostics Program and select the Configure Adapter menu item.

DISABLE SHADOW MEMORY BASE is the starting address to exclude. For the example above, it is the starting BIOS MMIO address.

The ISA IRQ setting must be set to USED if the adapter is configured for this IRQ. (IRQ 9 is the default setting for the TokenLink III ISA adapter.)

ODINSUP and Artisoft's LANTASTIC 5.0

As part of Novell's commitment to be interoperable, ODI allows users to support NDIS protocol stacks. An ODI module called ODINSUP.COM allows NDIS protocol stacks to run unmodified over the ODI LSL and talk to an ODI LAN driver.

To install the ODINSUP module, you simply install it in memory. In DOS, this is done by loading ODINSUP.COM either at the command line or in a batch file. (**Note:** The NDIS PROTMAN device driver must be loaded before the ODINSUP module can be loaded.)

To configure ODINSUP, you modify the AUTOEXEC.BAT file (or any batch file that calls on the ODINSUP file), as well as the CONFIG.SYS, NET.CFG, and PROTOCOL.INI files. If you are using LANTASTIC® 5.0, the PROTOCOL.INI file is replaced by the STARTNET.BAT file that BINDS the NDIS protocol stack to the ODI stack.

New Independent-Study Course

3Com's new independent-study course, High-Speed Data Networking (3CS-390), captures the basics of FDDI, FDDI II, 100BASE-T Fast Ethernet, 100VG-AnyLAN, ATM, SONET, and SMDS explaining why and when to implement each one. This forward-looking course provides enough technical detail for you to make informed decisions about high-speed features, limitations, and what lies ahead.

This leading-edge course is priced at $360. To order, call your reseller or call 3Com Education Services at (800) 876-3266, press option 7.

The following example files are configured for 3Com EtherLink III adapters; if you use a different network adapter, change the marked lines accordingly.

AUTOEXEC.BAT file

```
LSL
3C5X9.CO      Change to the ODI driver used by your adapter if needed.
IPXODI
ODINSUP
NETBIND
NETX
```

CONFIG.SYS file

```
DEVICE=C:\LANTASTI\PROTMAN.DOS /I:C:\LANTASTI
Do not include an NDIS device driver for the network adapter in this file.
LASTDRIVE=E      (Use "E" with NETX and "Z" with VLM)
```

NET.CFG file

```
PROTOCOL ODINSUP
     BIND 3C5X9      Change to the ODI driver used by your adapter if needed.

LINK DRIVER 3C5X9      Change to the ODI driver used by your adapter if needed.
     FRAME ETHERNET_802.3
     FRAME ETHERNET_802.2
     FRAME ETHERNET_II
     FRAME ETHERNET_SNAP
     PROTOCOL IPX 0 ETHERNET_802.3
```

STARTNET.BAT file

```
@ECHO OFF
PATH     (Specify the path needed.)
SHARE /L:200

AI-NDIS BIND_TO=X3C5X9      Change to the ODI driver used by your adapter if needed.

REDIR ASHES LOGINS=3      (Example only.)
SERVER
```

PROTOCOL.INI file

```
[PROTMAN]
DRIVERNAME=PROTMAN$
DYNAMIC = YES
PRIORITY = AILANBIO

;3COM 3C509
;DO NOT put any I/O, INTERRUPT etc. settings in this file for this
;adapter. The following two lines are all that are needed.

[ETHERLINKIII]
DRIVERNAME = ELNK3$
```

3Com Quick Technical Resource List

3Com World Wide Web

For Internet access to 3Com service features, news, product information, and job postings, drop by the web page at (URL) "http://www.3com.com/".

3Com BBS

3Com BBS, formerly CardBoard®, is a bulletin board system that contains software drivers, patches, fixes, technical tips, product information, diagnostic programs, and card services for the latest PCMCIA compatibility information. You can download data via modem, seven days a week, 24 hours a day.

How to Access 3Com BBS

To reach the service, set your modem to 8 data bits, no parity, and 1 stop bit. Then choose up to 14400 baud signaling rate and dial 1-408-980-8204 from Canada, Latin America, or the U.S. Outside the U.S., Canada, and Latin America, contact your local 3Com sales office.

Product Information

800-638-3266; outside the U.S. and Canada, call your local 3Com sales office

Technical Support

800-876-3266, press option 1; outside the U.S. and Canada, call your local 3Com sales office

Pinouts for NETBuilder Products

The following tables list the connector pinout assignments for the various serial interfaces available on 3Com's NETBuilder® bridge/router products.

DB-9 Console Port Pinouts

Pin Number	Signal Description
1	Data Carrier Detect
2	Receive Data
3	Transmit Data
4	Data Terminal Ready
5	Signal Ground
6	Data Set Ready
7	Request to Send
8	Clear to Send
9	Ring Indicator

RS-232 Connector Pin Assignments

Pin Number	RS-232 Name	Signal Description
1	AA	Chassis Ground
2	BA	Transmit Data
3	BB	Receive Data
4	CA	Request to Send
5	CB	Clear to Send
6	CC	Data Set Ready
7	AB	Signal Ground
8	CF	Data Carrier Detect
15	DB	Transmit Clock
17	DD	Receive Clock
20	CD	Data Terminal Ready
24	DA	External Transmit Clock

HSS V.35 Module V.35 and RS-232 Connector Pin Assignments

HSS V.35 Pin Number	Signal Name	V.35 Pin Number/ID	RS-232 Pin Number
1	Chassis Ground	A	1
2	TD		2
3	RD		3
4	RTS	C	4
5	CTS	D	5
6	DSR	E	6
7	Signal Ground	B	7
8	DCD	F	8
9	RDA	R	
10	RDB	T	
11	unassigned		
12	SCRA	V	
13	SCRB	X	
14	SDA	P	
15	TXCI		15
16	SDB	S	
17	RXCI		17
18	SCTEA	U	
19	SCTEB	W	
20	DTR	H	20
21	SCTA	Y	
22	SCTB	a	
23	RXCO		
24	TXCO		24
25	unassigned		

3Com Quick Technical Resource List (Continued)

Service Contract Help Line
800-876-3266, press option 3; outside the U.S. and Canada, call your local 3Com sales office

Training and Independent Study Courses
800-876-3266, press option 7, select 1
Fax: 408-764-7290
Outside the U.S. and Canada, call your local 3Com sales office

Ask3Com℠ Information Service on CompuServe®
800-848-8199; outside the U.S. and Canada, call the nearest CompuServe office

HDWAN V.35 Cable and 68-Pin Connector Pinout

Adapter A	Cable B	Pinout C	Signal	I/O	Signal Description	V.35 Pin
57	46	35	GND		Shield GND	A
28	17	6	SDB	output	TX data, non-inverted	S
29	18	40	SDA	output	TX data, inverted	P
32	20	8	RDB	input	RX data, non-inverted	T
66	54	42	RDA	input	RX data, inverted	R
58	47	36	RTS	output	Request to Send	C
60	49	37	CTS	input	Clear to Send	D
26	15	4	DSR	input	Data Set (DCE) Ready	E
24	13	1	DTR	output	Data Terminal (DTE) Ready	H
27	16	5	GND		Signal GND	B
61	50	38	DCD	input	Data Carrier Detect	F
34	22	10	SCTB	input	Transmit clock, non-inverted	AA
68	56	44	SCTA	input	Transmit clock, inverted	Y
33	21	9	SCRB	input	Receive clock, non-inverted	X
67	55	43	SCRA	input	Receive clock, inverted	V
25	14	2	LL	output	Local Loop Back	L
59	48	3	RL	output	Remote Loop Back	N
62	51	39	RI	input	Ring Indicator	J
30	19	7	SCTEB	output	Ext. Transmit clock, non-inverted	W
64	53	41	SCTEA	output	Ext. Transmit clock, inverted	U
23	12	11	TM	input	Test Mode	NN
65	52	45	GND		Shield GND	
31	63		NC		not connected	

RS449 Connector Pinout

Pin Number	Signal	Signal Description	Pin Number	Signal	Signal Description
1	Shield	Chassis Ground			
2	SI	Signal Rate Indicator	20	RC	Receive Common
3		spare	21		spare
4	SD (A)	Send Data	22	SD (B)	Send Data
5	ST (A)	Send Timing	23	ST (B)	Send Timing
6	RD (A)	Receive Data	24	RD (B)	Received Data
7	RS (A)	Request to Send	25	RS (B)	Request to Send
8	RT (A)	Receive Timing	26	RT (B)	Receive Timing
9	CS (A)	Clear to Send	27	CS (B)	Clear to Send
10	LL	Local Loopback	28	IS	Terminal In Service
11	DM (A)	Data Mode	29	DM (B)	Data Mode
12	TR (A)	Terminal Ready	30	TR (B)	Terminal Ready
13	RR (A)	Receiver Ready	31	RR (B)	Receiver Ready
14	RL	Remote Loopback	32	SS	Select Standby
15	IC	Incoming Call	33	SQ	Signal Quality
16	SF/SR+	Select Frequency/ Signal Rate Selector	34	NS	New Signal
17	TT (A)	Terminal Timing	35	TT (B)	Terminal Timing
18	TM	Test Mode	36	SB	Standby Indicator
19	SG	Signal Ground	37	SC	Send Common

RS449 37-Pin and X.21 15-Pin Adapter Cable Pinouts

RS449 Name	Pin	X.21 Pin	Name
Shield	1	1	Shield
SD (A)	4	2	TX Data A
SD (B)	22	9	TX Data B
RD (A)	6	4	RX Data A
RD (B)	24	11	RX Data B
ST (A)	5	6	Clock A
RT (A)	8		
ST (B)	23	13	Clock B
RT (B)	26		
TR (A)	12	3	Control A
RR (A)	13		
TR (B)	30	10	Control B
RR (B)	31		
DM (A)	11	5	Indicate A
DM (B)	29	12	Indicate B
SG	19	8	Zero volts

Setting the IRQ Mask for Card and Socket Services

If device drivers for Card and Socket Services are loaded on a laptop or mobile computer, the PCMCIA controller is assigned system resources from a pool allocated during Card and Socket Services installation. In some cases, the PCMCIA controller can allocate restricted resources such as Interrupt Request 6 (Floppy IRQ), causing resource conflicts.

The best way to prevent the PCMCIA controller from allocating restricted resources is to use the IRQ mask function. The mask can be set by a command line switch or by a menu-driven utility. Consult your manual for the proper method of using the mask function.

To use a command line switch, insert the switch into the DEVICE= line of your computer's CONFIG.SYS file. The examples that follow show the command line switches for three commonly used Card and Socket Services drivers. Consult your manual for the correct syntax and method of setting this function for your driver.

Example 1

Device driver: Intel 82365SL
CONFIG.SYS statement: DEVICE=C:\CANDSS\SS365SL.EXE /IIRM:9D00h
IRQs masked: 0, 1, 2, 3, 4, 5, 6, 7, 9, 13, 14
Switch: /IIRM
Default value: DEB8h (IRQs masked: 0, 1, 2, 6, 8, 13)
Table 1 shows how the hex value in this example is translated into IRQ values.

Table 1. *IRQ Values*

HEX			9				D				0				0			
BIN	1	0	0	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0
IRQ	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		

Automated Fax Service

3Com's interactive fax service, 3ComFacts℠, provides technical information on 3Com products 24 hours a day, seven days a week.

Call 3ComFacts using your touchtone telephone. International calling numbers are:

Hong Kong
852 537 5610

United Kingdom
44 1442 278279

United States
1 408 727 7021

Free local access is available within the following countries using the numbers below:

Australia
1 800 123 853

Denmark
800 17319

Finland
98 001 4444

France
05 90 81 58

Germany
0130 81 80 63

Italy
1678 99085

The Netherlands
06 0228049

Norway
80 01 1062

Sweden
020 792954

United Kingdom
0800 626403

Example 2:
Device driver: DataBook TCIC-2/N
CONFIG.SYS statement: DEVICE=C:\SSDB\SSDBOOK.EXE /IIRM:9D94h
IRQs masked: 0, 1, 3, 5, 6, 9, 13, 14
Switch: /IIRM
Default: 4CF8h (IRQs masked: 0, 1, 2, 8, 9, 12, 13, 15)

Example 3:
Device driver: Cirrus CL-PD67x0
CONFIG.SYS statement: DEVICE=C:\CARDSOFT\SSCIRRUS.EXE /IRM:4690h
IRQs masked: 0, 1, 2, 3, 5, 6, 8, 11, 12, 13, 15
Switch: /IRM
Default: DEB8h (IRQs masked: 0, 1, 2, 6, 8, 13)

3Com U.S./Canada Sales Offices

3Com Headquarters
Santa Clara
Phone: 800-NET-3Com
or 408-764-5000
Fax: 408-764-5001

3Com U.S.A.
California
Costa Mesa
Phone: 714-432-6588
Fax: 714-432-6574

Los Angeles (Downtown)
Phone: 213-612-7728
Fax: 213-426-2115

Los Angeles
(LAX Airport)
Phone: 310-348-8110
Fax: 310-348-8167

Pleasanton
Phone: 510-847-2040
Fax: 510-463-1560

San Diego
Phone: 619-546-4476
Fax: 619-453-2839

San Francisco
Phone: 415-955-2607
Fax: 415-397-6309

Santa Clara
Phone: 408-764-5000
Fax: 408-764-6740

Colorado
Englewood
Phone: 303-694-0674
Fax: 303-694-1670

Florida
Boca Raton
Phone: 407-392-0668
Fax: 407-362-4375

Maitland (Orlando area)
Phone: 407-661-1999
Fax: 407-660-0012

Tampa
Phone: 813-281-4635
Fax: 813-281-4619

Georgia
Atlanta
Phone: 404-395-2370
Fax: 404-395-2378

Illinois
Itasca
Phone: 708-250-5400
Fax: 708-250-5407

Iowa
Des Moines
Phone: 515-245-3794
Fax: 515-245-3793

Massachusetts
Waltham
Phone: 617-466-9700
Fax: 617-890-9621

Michigan
Southfield
Phone: 810-262-1580
Fax: 810-262-1588

Minnesota
Minneapolis
Phone: 612-921-8335
Fax: 612-921-8327

Missouri
St. Louis
Phone: 314-984-6800
Fax: 314-984-6803

New Jersey
Cherry Hill
Phone: 609-751-4125
Fax: 609-424-7112

Rutherford
Phone: 201-531-9177
Fax: 201-531-9196

New York
New York City
Phone: 212-643-1733
Fax: 212-643-1738

North Carolina
Raleigh
Phone: 919-676-5292
Fax: 919-676-5294

Ohio
Cincinnati
Phone: 513-563-3559
Fax: 513-563-3523

Worthington
Phone: 614-438-4165
Fax: 614-438-4166

Pennsylvania
Pittsburgh
Phone: 412-928-4978
Fax: 412-928-4977

West Conshohocken
Phone: 215-941-2777
Fax: 215-825-6296

Tennessee
Brentwood
Phone: 615-377-0754
Fax: 615-661-5587

Texas
Dallas
Phone: 214-980-5005
Fax: 214-980-5020

Houston
Phone: 713-864-3399
Fax: 713-864-4098

Utah
Salt Lake City
Phone: 801-264-6690
Fax: 801-264-6691

Virginia
Richmond
Phone: 804-273-0488
Fax: 804-273-0791

Vienna
Phone: 703-749-4200
Fax: 703-749-7988

Washington
Bellevue
Phone: 206-455-8530
Fax: 206-454-7402

3Com CANADA
Alberta
Calgary
Phone: 403-265-3266
Fax: 403-265-3268

British Columbia
Vancouver
Phone: 604-434-3266
Fax: 604-434-3264

Ontario
Ottawa
Phone: 613-782-2901
Fax: 613-782-2391

Toronto
Phone: 416-498-3Com
Fax: 416-498-1262

Quebec
Montreal
Phone: 514-874-8008
Fax: 514-393-1249

3Com Sales Offices Worldwide

To reach a 3Com office outside the country you are calling from, use the country code shown in parentheses and precede the number with the international access code of the country you are calling from (for example, 010 from the United Kingdom).

3Com HEADQUARTERS
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, California 95052-8145
Phone: (1) 408-764-5000
Fax: (1) 408-764-5001

3Com Limited
Hemel Hempstead, England
Phone: (44) 442 278000
Fax: (44) 442 278003

3Com Europe
France, Israel, North Africa
3Com France
Paris, France
Phone: (33) 1 69 86 68 00
Fax: (33) 1 69 07 11 54

Austria, Bulgaria, Czech Republic, Germany, Hungary, Poland, Romania, Slovakia, Switzerland
3Com GmbH
Munich, Germany
Phone: (49) 89 62732 0
Fax: (49) 89 62732 233

Albania, Greece, Italy, Malta
3Com Mediterraneo SRL
Milan, Italy
Phone: (39) 2 273 02041
Fax: (39) 2 273 04244

Portugal, Spain
Madrid, Spain
Phone: (34) 1 383 1700
Fax: (34) 1 383 1703

The Netherlands
3Com Benelux, B.V.
Nieuwegein, The Netherlands
Phone: (31) 3402 55033
Fax: (31) 3402 54630

Belgium, Luxembourg
Diegem, Brussels
Phone: (32) 2 716 4880
Fax: (32) 2 716 4780

Baltic States, Denmark, Finland, Iceland, Norway, Sweden
3Com Nordic AB
Stockholm, Sweden
Phone: (46) 8 632 9100
Fax: (46) 8 632 0905

England, Russia, Scotland, Wales
3Com (U.K.) Limited
Buckinghamshire, U.K.
Phone: (44) 628 897000
Fax: (44) 628 897003

Ireland
3Com Ireland Limited
Dublin
Phone: 353 1 820 7077
Fax: 353 1 820 7107

South Africa
3Com South Africa
South Africa
Phone: 27 11 803 7404
Fax: 27 11 803 7411

Bahrain, Jordan, Kuwait, Lebanon, Oman, Qatar, Saudi Arabia, Syria, United Arab Emirates, Yemen
3Com Middle East
Dubai, United Arab Emirates
Phone: 971 4 349 049
Fax: 971 4 349 803

3Com Asia Pacific Rim
Australia, New Zealand
3Com ANZA East
North Sydney, Australia
Phone: 61 2 959 3020
Fax: 61 2 956 6247

3Com ANZA West
Melbourne, Australia
Phone: 61 3 653 9515
Fax: 61 3 653 9505

Hong Kong
3Com Asia Ltd
Hong Kong
Phone: 852 2501 1111
Fax: 852 2537 1149

Indonesia
3Com Asia Ltd
Phone: 6211 523-9181
Fax: 6211 523-9156

Japan
3Com Japan, K.K., Tokyo
Phone: 81 3 3345 7251
Fax: 81 3 3345 7261

Korea
3Com Asia Ltd
Seoul
Phone: 822 732 4434
Fax: 822 732 4437

Malaysia
Kuala Lumpur
Phone: 60 3 233 6162
Fax: 60 3 233 6174

People's Republic of China
3Com Asia Ltd
Beijing
Phone: 861 849 1380
Fax: 861 849 1381

Singapore
3Com Asia Ltd
Phone: 65 538 9368
Fax: 65 538 9369

Taiwan
Taipei
Phone: 886 2 377-5850
Fax: 886 2 377-5860

3Com Latin America
Latin America
3Com Corporation
Santa Clara, California
Phone: (1) 408-764-5730
Fax: (1) 408-764-5742

Brazil
3Com do Brazil
Sao Paulo-SP, Brazil
Phone: 55 11 530 2318
Fax: 55 11 241 1571

Chilé
3Com de Chilé
Santiago
Phone: 562 633-9242
Fax: 562 633-4338

México
3Com de México
Col. Granada, México
Phone: (011) 525-531-0591
Fax: (011) 525-254-3159

Miami
3Com Latin America
Miami, Florida
Phone: (1) 305-261-3266
Fax: (1) 305-261-4901

Calendar

Upcoming Events

April 19–22
PC World/ Networking China 1995
Shanghai, China

April 20
NUI – Orlando
Orlando, Florida

May 2–4
ATUG '95
Sydney, Australia

May 17
NUI – Vancouver
Vancouver, Canada

May 17–19
Network Management '95
Paris, France

May 30–June 1
Networks '95 Copenhagen
Copenhagen, Denmark

For information on 3Com seminars in your area, call your local 3Com sales office.

3TECH Subscription Information

3TECH Subscription Information

❑ **Change of Address**
I already receive *3TECH*, but please make the following corrections.

❑ **Subscription Cancellation**
Please cancel my *3TECH* subscription.

PLEASE PRINT:

Name

Title

Company

Mailing address (street number or P. O. box number)

City/state or province | Zip+extension or postal code | Country

Country/area code | Telephone number | Country/area code | Fax number

Inside and Outside U.S./Canada

Mail this form to:
3Com Corporation
Attn: Linda Webb
P.O. Box 58145
Santa Clara, CA 95052-8145

Or fax this form to:
408-764-6477
Attn: Linda Webb

3Com®

ADDRESS CORRECTION REQUESTED

Bulk Rate
U.S. Postage
PAID
Sunnyvale, CA
Permit No. 5

3Com Corporation
P.O. Box 869038, Plano, Texas 75086
Printed in USA 740000-006