




Layer 3 Switching

An Introduction



Layer 3 Switching

An Introduction

Contents

A Brief History of Protocol Layering	2
Effective Information Management	2
Layering 101	3
Contemporary Layering Model	3
Layer 1	4
Layer 2	4
Layer 3	5
Layer 4	6
Layer 5	6
Enter the Layer 3 Switch	6
A Layer 3 Switch Is a Router	7
Router Interfaces as Layer 2 Switching Domains	7
Effective Application of Policy	7
Ease of Management	8
Layer 3 Switching v. Traditional Routing	8
The Business Case for a Layer 3 Switch	8
Cost Savings	10
Case Study	10
Initial Network	11
Phase 1	11
Phase 2	12
Phase 3	12
Future Trends	12
Conclusion	13

Acronyms and Abbreviations

ASIC

application-specific integrated circuit

ATM

Asynchronous Transfer Mode

BGP

Border Gateway Protocol

CoS

Class of Service

DHCP

Dynamic Host Configuration Protocol

DNS

Domain Name System

FDDI

Fiber Distributed Data Interface

FIRE

Flexible Intelligent Routing Engine

IEEE

Institute of Electrical and Electronics Engineers

IETF

Internet Engineering Task Force

IP

Internet Protocol

IPsec

IP Security

ISO

International Standards Organization

ISP

Internet service provider

MAC

Media Access Control

MIB

Management Information Base

MPOA

Multiprotocol over ATM

Layer 3 Switching

An Introduction

By Robert Ciampa

Many people involved in the deployment of information technology compare their profession to the world of Indiana Jones, a Hollywood action hero of great intelligence, challenged by friendly and unfriendly forces while searching the world for lost cities and hidden treasure. Like Indiana Jones, IT managers might not have a second chance if they make a wrong decision. In fact, the journey through the networking jungle is full of deception, wrong turns, and stumbling blocks.

In a competitive environment, the proper technology decisions can catapult corporations over their competitors, erasing barriers to entry and redrawing the battlefield. A prime example is amazon.com, which has used the Internet to revolutionize the bookselling industry, at the expense of formidable competitors who now attempt to mimic the techniques of their upstart foe. On the other hand, blindly following temporary technology fashion may leave IT managers stuck in the jungle, or out of a job.

This paper is a map through the jungle of internetworking infrastructure, particularly focusing on how Layer 2 switching and Layer 3 routing have combined to form the powerful Layer 3 switching architecture. The paper analyzes Layer 3 switching from both a functional and an operational perspective, helping the reader make an informed assessment of its merits as an enabling technology.

A Brief History of Protocol Layering

To fully appreciate Layer 3 switching, it is useful to examine its ancestry, since many common traits still prevail. Rather than go back to the stone age of hierarchical networking, we'll begin with the "modern era" of data communications, a time of peer-to-peer networking with heterogeneous systems. It is interesting to note that hierarchical networking—its best example being IBM's Systems Network Architecture (SNA)—was probably the quintessential—but immutable—client/server

architecture. SNA's formal counterpart, the International Standards Organization (ISO) Open Systems Interconnect (OSI) model, which was a seven-element layout, succeeded more as a pedagogical tool than as an implementation foundation. As a result, many academics, along with the some rare implementers (Digital Equipment Corporation with DECnet Phase V) were left in the networking jungle.

Meanwhile, the Internet Protocol (IP) was enjoying some deployment success, first through the U.S. Department of Defense's ARPANET—the genesis of the Internet—then into diverse university communities. IP and its associated higher-layer protocols, such as User Datagram Protocol (UDP) and Transport Control Protocol (TCP), were supposed to be supplanted by the OSI protocols, but the increasing complexity of OSI, exacerbated by a prolonged ratification process, undermined its prospects. IP continued to be deployed, while other IP-like protocols such as NetWare's IPX and Apple's AppleTalk were enjoying their own success. The similarity among IP, IPX, and AppleTalk is no accident: they share a common lineage through Xerox Network Systems (XNS), an older but simpler model than OSI.

Effective Information Management

Just as there are many types of jungles, so there are many types of data networks. And jungles and networks have some striking similarities in the way they are organized. In the jungle, the parts of the whole are called ecosystems; in the network, they are called layers. Each subsystem, or layer, is often quite distinct from others within the same system or network, but depends upon access to the others for its survival. Call it the food chain or call it the protocol stack.

Knowledge of layering is crucial for the strategic and tactical deployment of both networking and information technology in an organization. Many people view layering as an academic exercise in which Layer 2 represents switching and Layer 3 represents routing. Such shortsighted thinking leaves many organizations at the mercy of the performance

constraints of their collapsed backbone routers. Understanding the capabilities and limits of each layer is the foundation for information management. Strategic decisions must be made about application deployment, network scalability, performance, and cost of ownership. Tactical decisions must be made about which products to apply as part of an overall solution. This methodology becomes even more important as voice, video, and data networks continue to converge, blurring the once clear demarcation between data communications and telecommunications.

Layering 101

Although this paper is about Layer 3 switching, a quick overview of layering is needed. Layering schemes provide guidelines, rather than strict rules, for delegating networking functionality. Figure 1 shows the basic principles of layering. Elements at the same layer, shown on the horizontal, are known as peers and communicate via a well-known (and documented) protocol. Messages are exchanged among peers, the protocol defining the format, syntax, semantics, and sequencing. Elements within the same stack, shown on the vertical, communicate via an internal interface. This interface, though usually not well documented nor standard, often exhibits the same characteristics as a protocol, the only difference being that the interface protocol between Layer n and Layer $n + 1$ on stack 1 may be wholly different from that of stack 2.

As mentioned, communication within one stack may be different from that within




Figure 1. Layering Reference Model

Application
Presentation
Session
Transport
Network
Data link
Physical

Figure 2. OSI Layering Model

other stacks and entirely proprietary, but communication between peers in different stacks must be open and consistent. The notion of open systems has been a major factor in the growth and operation of the Internet, along with those of institutional organizations. It is also important to note that an element at a particular layer may be further broken down into additional layers. This is most clearly seen with Asynchronous Transfer Mode (ATM) models. Finally, in certain models, higher layers may share information with lower layers to either conserve system resources or improve performance. The Internet Engineering Task Force (IETF) Next-Hop Resolution Protocol (NHRP) is an example of this intra-layer communication, allowing Layer 3 “shortcuts.” This concept will be discussed later.

Contemporary Layering Model

For many years, the OSI model (Figure 2) was the reference layering paradigm for data networking. The OSI model was an extremely powerful architecture that included well-defined Layer n /Layer $n + 1$ protocols in addition to rich peer-to-peer protocols. Unfortunately, much of this model succumbed to the complexity of the protocols and the effects of an overly rigorous standardization process. Since only a few elements survived to become part of the contemporary networking model, no further analysis will be made of this model.

The contemporary network layer architecture is much simpler than its OSI counterpart. Originating from various research and defense initiatives, the contemporary model was intended to be supplanted by OSI. Instead, it became the de facto networking

Acronyms and Abbreviations

NHRP

Next-Hop Resolution Protocol

OSI

Open Systems Interconnect

OSPF

Open Shortest Path First

QoS

Quality of Service

RIP

Routing Information Protocol

RMON

Remote Monitoring

RSVP

Resource Reservation Protocol

RTP

Real Time Protocol

SDH

Synchronous Digital Hierarchy

SLA

Service Level Agreement

SNA

Systems Network Architecture

SONET

Synchronous Optical Network

TCP

Transport Control Protocol

UDP

User Datagram Protocol

VPN

virtual private network

WAN

wide area network

WinSock

Windows Sockets

XNS

Xerox Network Systems

Application
Transport
Routing
Switching
Interface

Figure 3. *Contemporary Layering Model*

standard, especially through IP. As mentioned, both IPX and AppleTalk are quite similar to IP, but are slowly becoming less prominent as IP dominance continues to grow. This discussion will emphasize IP, but the methods discussed can easily be applied to environments using NetWare and Apple protocols.

Figure 3 shows the contemporary networking model based upon IP. Network participants, whether infrastructure equipment (switches and routers) or end systems (clients and servers), may include some or all of the protocol stack.

Layer 1

This layer, known as the interface layer, is responsible for device connectivity. Though usually represented by well-known network types—Ethernet, Fast Ethernet, Gigabit Ethernet, Token Ring, FDDI, ATM, SONET/SDH, etc.—Layer 1 also covers the subtypes. For example, Fast Ethernet provides physical connectivity over copper media (100BASE-TX) and over fiber media (100BASE-FX). Fiber can be further divided into multimode or single mode, with single mode further partitioned based on its “reach,” the distance over which it can transmit. Certain technologies are actually used as a pure Layer 1 element (SONET/SDH) or provide a virtual Layer 1 element (ATM with SONET/SDH).

While the various types of Ethernet are rather straightforward, FDDI, ATM, and SONET/SDH add more complexity, while providing extended Layer 1 capabilities such as fault tolerance and support for physical multiplexing to support distinct traffic flows such as voice and data. With these added capabilities comes added cost, and sometimes slower performance.

Layer 2

This layer, known as the switching layer, allows end station addressing and attachment. Because architectures up to Layer 2 allow end station connectivity, it is often practical to construct a Layer 2-only network, providing simple, inexpensive, high-performance connectivity for hundreds or even thousands of end stations. The past five years have seen the extraordinary success of the “flat” network topologies provided by Layer 2 switches connected to other Layer 2 switches or ATM switches.

Layer 2 switching, also called bridging, forwards packets based on the unique Media Access Control (MAC) address of each end station. Data packets consist of both infrastructure content, such as MAC addresses and other information, and end-user content. At Layer 2, generally no modification is required to packet infrastructure content when going between like Layer 1 interfaces, like Ethernet to Fast Ethernet. However, minor changes to infrastructure content—not end-user data content—may occur when bridging between unlike types such as FDDI and Ethernet. Either way, processing impact is minimal and so is configuration complexity.

Layer 2 deployment has seen the most striking infrastructure change over the past decade. Shared Ethernet, represented by particular cable types or contained within shared hubs, offered a very simple, and even more inexpensive, approach for Layer 2. Though still quite popular, shared technology, where all stations use the same bandwidth slice, has very limited scaling capabilities. Depending upon the applications being used, shared networks of more than one hundred users are becoming less common. Many network designers have “tiered” their infrastructure by feeding shared Layer 2 into switched Layer 2 or even Layer 3. Switched Layer 3 apportions each station—or port—its own dedicated bandwidth segment. Recent enhancements at Layer 2 provide packet prioritization capabilities for the application of network policies. The new IEEE 802.1p standard defines Class of Service (CoS) policies capabilities for Layer 2 segments.

Note that Layer 2 does not ordinarily extend beyond the corporate boundary. To connect to the Internet usually requires a router; in other words, scaling a Layer 2 network requires Layer 3 capabilities.

Layer 3

This layer, known as the routing layer, provides logical partitioning of subnetworks, scalability, security, and Quality of Service (QoS). QoS, a recent enhancement to Layer 3, goes beyond the simple packet prioritization found in CoS by providing bandwidth reservation and packet delay bounding.

The backbone of the Internet, along with those of many large organizations, is built upon a Layer 3 foundation. IP is the premier Layer 3 protocol. In addition to Layer 2 MAC addresses, each IP packet also contains source and destination IP addresses. For an intranet packet, one IP address addresses the client, the other the server.

IP in itself is not a particularly complex protocol; extensive capabilities are supplied by the other components of the IP suite. The Domain Name System (DNS) removes the burden of remembering IP addresses by associating them with real names. The Dynamic Host Configuration Protocol (DHCP) eases the administration of IP addresses and is used extensively by network administrators and Internet service providers (ISPs). Routing protocols such as Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP) provide information for Layer 3 devices to direct data traffic to the intended destination. IP Security (IPsec) furnishes elements necessary for security, such as authentication and encryption. IP not only allows for user-to-user communication, but also for efficient dissemination over point-to-multipoint flows, known as IP Multicast. Higher-layer protocols, discussed later in this paper, provide even greater versatility for content distribution.

Although many organizations received tremendous performance advantages by converting routed and shared networks to Layer 2 switching, it became apparent that some level of partitioning was still required. Consequently,

routers maintained a presence at many points within a corporate network. For a while this presented minimal problems, since a majority of the data traffic stayed local to the subnet, which was increasingly being serviced by a Layer 2 switch. But concurrent with the increasing acceptance of Layer 2 switching as an essential component of network infrastructure were two other developments: the migration of servers to server farms for increased security and management of data resources; and the deployment of intranets, organization-wide client/server communications based on Web technology. These factors began moving data flows off local subnets and onto the routed network, where the limitations of router performance increasingly led to bottlenecks.

With the routers causing information flow constriction, IT managers became increasingly reluctant to deploy new, enabling technologies, such as multicast-based applications and middleware. Even the migration of desktops to higher-performance media connections, such as 100 Mbps Fast Ethernet, were scrutinized as long as 10 Mbps router interface funnels were in place.

Router vendors attempted to respond by offering higher-performance interface cards, but throughput was fundamentally bounded by centralized, software-based architectures that simply could not go any faster. The same software responsible for managing WAN links, X.25, and asynchronous terminal lines was now expected to handle next-generation gigabit networks. Router vendors tried distributing functionality to improve performance, resulting in a hodgepodge collection of route processing and interface cards. Was the device still routing, or was it performing some other packet forwarding scheme?

Emerging QoS was even more suspect. The IETF was moving forward on Resource Reservation Protocol (RSVP), a signaling method to set up bandwidth and delay control on packet-based internetworks. Monitoring RSVP flows, using a process known as policing, required extensive software support on already overburdened legacy routers. Could QoS be practical on a contemporary LAN?

Meanwhile, standards bodies such as the ATM Forum were working on methods to offload the Layer 3 bottleneck by exploiting the capabilities of the lower layers. One result was the Multiprotocol over ATM (MPOA) specification, which uses Layer 3 routing information and the IETF's NHRP protocol to offload the routers and provide forwarding at the physical (ATM) layer. A Layer 3 switch can route at Layer 3 or utilize MPOA; the performance is identical.

Layer 4

This layer, known as the transport layer, is the communication path between user applications and the network infrastructure and defines the method of communicating. TCP and UDP are well-known examples of elements at the transport layer. TCP is a "connection-oriented" protocol, requiring the establishment of parameters for transmission prior to the exchange of data. Web technology is based on TCP. UDP is "connectionless" and requires no connection setup, which is especially important for multicast flows. Elements at this level also differ in the amount of error recovery provided and whether or not it is visible to the user application. Both TCP and UDP are layered on IP, which has minimal error recovery and detection mechanisms, leaving the burden at Layer 4 or higher. TCP forces retransmission of data that was lost by the lower layers, whereas UDP makes the application responsible.

A major enhancement to multimedia support at Layer 4 is the Real Time Protocol (RTP). RTP works in conjunction with UDP, and provides services necessary for packet timing and sequencing. Many time-sensitive applications running over IP networks now actually include both UDP and RTP.

Layer 5

This layer, known as the application layer, provides access to either the end user or some type of information repository such as a database or data warehouse. Users communicate with the application, which in turn delivers data to the transport layer. Applications do not usually communicate with the lower layers;

rather, they are written to interface with a specific communication library, like the popular WinSock library available in Windows-based workstations.

When defining the behavior of the applications they are writing, developers decide on the type of transport mechanism necessary. For example, database or Web access requires robust, error-free access and would demand TCP, though it could be implemented with more code and in a more cumbersome manner with UDP. Multimedia, on the other hand, cannot tolerate the overhead of connection-oriented traffic and will commonly make use of UDP. For prioritization, either TCP nor UDP can be selected, depending on the application or other parameters such as time of day. Any assistance that a network device can provide in terms of prioritization of the application would be extremely beneficial to the network manager, particularly during times of traffic volume from the LAN to the WAN.

Enter the Layer 3 Switch

Traditional routers, once the core components of enterprise networks, became a major obstacle to the migration toward next-generation networks. All the magic and alchemy involved in trying to make a software-based router forward packets more quickly created only an illusion of serious Layer 3 routing performance. A fundamental shift in technology was required.

In 1992, 3Com, a pioneer in both Layer 2 switching and traditional routing, began integrating its switching and routing products. The motivation was twofold: to reduce the number of devices to be managed, and to lower the cost of a combined Layer 2 and 3 solution. Though the first solution was mostly software based, subsequent products displayed increasing use of application specific integrated circuits (ASICs)—first for address table queries, then for forwarding packets. In 1997, 3Com delivered its third-generation, fully ASIC-enabled Layer 3 CoreBuilder™ 3500 switch, based on the Flexible Intelligent Routing Engine (FIRE) ASIC. For more information on FIRE, see the white paper "Flexible Intelligent Routing Engine (FIRE): The Third-

Table 1. Layer 3 Switching Product Evolution

Generation	Technology	Product	Routing Performance
First	Software	LANplex® 5000 switch	50K pps
Second	ISE ASIC	CoreBuilder 2500, 6000 switch	100K–1.1M pps
Third	FIRE ASIC	CoreBuilder 3500, 9000 switch	3.5M–64M pps

Generation Layer 3 Switching Architecture from 3Com,” available at www.3com.com.

Table 1 shows 3Com’s Layer 3 switching product evolution.

A Layer 3 Switch Is a Router

Vendors and the trade press alike have tried to apply the term “Layer 3 switch” to various products of the day, succeeding only in confusing IT decision makers. This paper aims to remove that confusion. A Layer 3 switch does everything to a packet that a traditional router does:

- Determines forwarding path based on Layer 3 information
- Validates the integrity of the Layer 3 header via checksum
- Verifies packet expiration and updates accordingly
- Processes and responds to any option information
- Updates forwarding statistics in the Management Information Base (MIB)
- Applies security controls if required

Because it is designed to handle high-performance LAN traffic, a Layer 3 switch can be placed anywhere within a network core or backbone, easily and cost-effectively replacing the traditional collapsed backbone router. The Layer 3 switch communicates with the WAN router using industry-standard routing protocols like RIP and OSPF.

Router Interfaces as Layer 2 Switching Domains

The Layer 3 switch has inherent Layer 2 switching domains per interface, allowing for individual subnet bandwidth allocation, along with broadcast containment. Not all interfaces are created equal, so the ability to group ports

together, whether based on physical characteristics or protocol information, is an extremely powerful tool for network designers concerned with capacity planning. This architecture is inherently scalable, capable of supporting numerous external Layer 2 switches that reside either in the data center or the wiring closet.

Such a design model preserves the subnetted infrastructure, concurrently boosting performance of those subnets and enabling the deployment of switched 10, 100, or 1000 Mbps right to the desktop if so desired. The concept of “subnet preservation” is the key to effective and trouble-free network migration—it allows gradual migration, helping IT managers to work within their staffing constraints without the need to renumber and reassign their entire network.

Effective Application of Policy

As previously stated, contemporary Layer 3 switches perform their forwarding—whether Layer 2, Layer 3, unicast, multicast, or broadcast—in hardware. Software is deployed to handle network administration, table management, and exception conditions. Some technologists view the hardware component of a Layer 3 switch as inflexible. In fact, hardware provides the ultimate flexibility not only in performance, but in parallel processing as well. The parallel processing model allows the network device to perform far more operations on packets than previously imagined, especially with respect to the application of policy.

A policy is a mechanism to alter the normal forwarding of a packet through a networking device. Familiar examples include security, load balancing, and protocol option

processing. Newer policies include QoS, a way to allocate bandwidth and control propagation delay, in addition to CoS, a way to manage packet prioritization. QoS and CoS policies are not only meant to enable new multimedia applications, such as LAN telephony, but to ensure network response time for mission-critical applications, such as telemedicine. Policy implemented by intelligent networking devices, such as Layer 3 switches, enables the integration of voice, video, and data onto the same infrastructure, a process 3Com calls convergence.

Software-based architectures cannot seamlessly administer policy controls at even moderate rates of speed (beyond 10 Mbps). The Layer 3 switch solves the problem, enabling policies to be applied at the same performance levels as ordinary Layer 2 and 3 forwarding. Further innovation allows the Layer 3 switch to apply policy based on Layer 4 information, such as TCP and UDP port information. Forward thinkers refer to this as "Layer 4 switching." The FIRE architecture supports all these policies, all the way to Layer 4.

Even with the massive capacity additions being planned for many networks, effective policy management enabled by Layer 3 switching is key to the protection and availability of critical resources.

Ease of Management

One of the critical success factors for the Layer 2 switch was its implementation and operational simplicity. Deployment was often as easy as powering on the switch, assigning it an IP address, and making the physical network connections. Routers, on the other hand, required extensive training and forced users to sift through a multitude of arcane commands. Layer 3 switches remove such complexity. Setting up a routed environment is as simple as setting up a Layer 2 switch, defining the routed interface, and enabling the routing protocols. IT managers concerned about their investment in training staff on traditional router platforms must assess whether this is truly an investment, or simply a sunk cost based upon vendor lock-in schemes.

For the network management application perspective, a Layer 3 switch behaves exactly as a legacy router does. Because of its Layer 2 component, extensive Remote Monitoring (RMON) capabilities are available. However, since Layer 3 and Layer 4 capabilities are present in the Layer 3 switch, higher-layer monitoring is available with RMON2 technology. RMON and RMON2 have historically been deployed with expensive external devices known as probes. Moving the RMON/RMON2 capability into the Layer 3 switch is a major benefit for IT administrators.

Layer 3 Switching vs. Traditional Routing

By now, it should be clear that a Layer 3 switch can be deployed anywhere in the LAN where a traditional router can be or has been used.

Table 2 compares the two types of devices. The Layer 3 switch has been optimized for high-performance LAN support and is not meant to service wide area connections (although it could easily satisfy the requirements for high-performance MAN connectivity, such as SONET). This optimization boosts the performance of a Layer 3 switch to as much as ten times that of a legacy router, while driving the price down to as little as a tenth. This cost comparison does not include the lower training costs for Layer 3 switch administrators or the increased productivity of a high-performance network.

There is another major architectural difference between a Layer 3 switch and a router. A traditional router organizes bridging (Layer 2) and routing (Layer 3) as peers. A Layer 3 switch layers routing on top of switching, permitting a more natural networking architecture while greatly facilitating scalability.

The Business Case for a Layer 3 Switch

Some IT managers may be concerned about deploying a "new" technology such as Layer 3 switching to their network. But Layer 3 switching is really an integration of two proven technologies: switching and routing. In fact, some Layer 3 switches are running the exact same routing software that has been fully tested and used in mission-critical networks

Table 2. *Layer 3 Switch vs. Legacy Router*

Characteristic	Layer 3 Switch	Legacy Router
Routes core LAN protocols: IP, IPX, AppleTalk	Yes	Yes
Subnet definition	Layer 2 switch domain	Port
Forwarding architecture	Hardware	Software
RMON support	Yes	No
Price	Low	High
Forwarding performance	High	Low
Policy performance	High	Low
WAN support	No	Yes

for nearly a decade. So whether the decision maker is an early adopter of technology or more conservative, the Layer 3 switch can satisfy both needs.

The first step toward the deployment of next-generation IT infrastructures is to ignore the networking element. Although this may seem absurd, it allows managers to focus on the end users, services, and data without being bound by historical network deficiencies. The network should be transparent. When the requirements for information transfer are known, capacity planning techniques will determine the necessary client and server interconnects. Organizational and security mandates are then applied, yielding the policy and subnetted infrastructure. Cost is then factored in. Finally, the decision is made regarding the appropriate networking products to satisfy these requirements.

Layer 3 switching technology must be considered from two perspectives. First, as a migration tool to move users forward to higher-performance networking, or surprisingly, to squeeze more performance out of what is currently installed. Many users complain about FDDI performance, only to discover that the network is running at less than 20 percent of capacity. The problem is not the network, but rather the devices attached to it. The second perspective addresses what can be done when network performance bottlenecks are removed. A high-performance network

enables a variety of steps to reduce costs and enhance security and business operations. The following are examples of several such steps.

- **Server farms.** Today, the viability of many organizations is closely related to their intellectual property, often stored on databases or server devices. The security and protection of these servers has been a major goal of IT managers, who have been at odds with the users of those servers. The point of contention has been the dependence of server performance on the network topology. The response has been to move servers within the same subnet or Layer 2 switching domain as users. With data traffic patterns becoming more distributed, this approach was breaking down. The Layer 3 switch allows servers to be centralized with no performance penalty, eliminating the cost of numerous server repositories while keeping end users satisfied.
- **Intranets.** Because of its secure nature, along with its higher capacity, the intranet is becoming a viable corporate communications vehicle with usage that includes HR record retrieval, major announcements, computer-based training, and live video broadcasts. Delivering a wide variety of services, some requiring a huge amount of bandwidth, can wreak havoc on the old router. The Layer 3 switch, because of its higher performance, traffic prioritization,

and subnet preservation, is ideally suited for the deployment of intranets.

- **Converged networks.** For some time, technological prognosticators have been extolling the virtues of multimedia and warning of the excessive demand it will place on IT infrastructures. But many network managers have been disinclined to integrate their voice, video, and data traffic, concerned not only with the bandwidth requirements, but fearing the degraded quality of the respective elements. The ability to recognize and respond to the unique attributes of voice, data, and video not only makes their integration viable, but also attractive from a cost and management perspective. The inherent flow recognition capabilities of Layer 3 switching enable practical deployment of converged networks without performance uncertainties.

Cost Savings

A traditional router may run U.S. \$8,000 to \$10,000 per Fast Ethernet interface, while a Layer 3 switch costs less than U.S. \$1,000 per

port. Surprisingly, greater densities can be achieved with Layer 3 switching, freeing up valuable rack space and saving on physical cabling plant expansion. Training costs plummet, too, as a four- to seven-day legacy router course is replaced with a one- or two-day class for the Layer 3 switch.

Major savings also lie in the ancillary effects of applying Layer 3 switching technology. Cost savings realized through server centralization, notably in physical plant and security, can be substantial, especially when space is at a premium. Other, less tangible effects include improved response time and conformance with SLAs. Clearly, the overall cost of ownership benefits of Layer 3 switches versus routers can be substantial.

Case Study

The following application scenario starts with a common contemporary network topology and illustrates a migration path toward a next-generation infrastructure. The deployment objectives are as follows:

- To minimize network disruption




Figure 4. Initial Network Configuration




Figure 5. Data Center and Workgroup Upgrades to Layer 3 Switching

- To preserve subnet infrastructure
- To avoid parallel network construction

Initial Network

The network core, shown in Figure 4, consists of an FDDI backbone running at 20 percent capacity, occasionally peaking at 40 percent. Collapsed backbone routers are the connection points to the FDDI backbone, with the exception of some data center servers that attach directly to the backbone. The legacy routers supply mostly 10 Mbps Ethernet interfaces, with some 100 Mbps Fast Ethernet interfaces. Some of these Ethernet interfaces connect to Layer 2 switches, which then cascade to hubs, while others connect to hubs directly. A majority of the desktops are shared 10 Mbps Ethernet. Some of the servers are switched. Departments may have their own server co-located on a subnet. The network is running IP and IPX, with the subnets for both protocols aligned with the other. The FDDI ring contains one subnet for each protocol, and each router interface also provides a sub-

net for each protocol. Two of the routers service WAN access: one for corporate network extension, the other for Internet service.

Phase 1

The first phase (Figure 5) consists of key legacy router replacement for the data center and for the most heavily used departments. If other legacy protocols such as DECnet or Banyan VINES are present, the Layer 3 switch and the router can be co-located in the wiring closet or the data center, the Layer 3 switch becoming the “express lane” for the contemporary protocols. With the Layer 3 switch in place, the department and data infrastructure behind it can then be upgraded to higher-performance Layer 2 switches, ultimately bringing switched Ethernet to the desktop. The migration of key departmental servers to server farms may begin at this point. If the capacity required for the aggregate client-server flows exceeds that of FDDI, the gradual evolution of the backbone may begin at this point, otherwise it will be covered in phase 2.




Figure 6. Backbone Migration Begins

Phase 2

The second phase (Figure 6 on page 12) continues the replacement of the routers on the FDDI backbone with Layer 3 switches. The routers that were servicing the WAN connections remain, but are now removed from the backbone and connected via Ethernet or Fast Ethernet to a Layer 3 switch. The migration of the backbone begins at this stage with the choice of Gigabit Ethernet or ATM. (This choice depends upon a variety of factors, which are beyond the scope of this paper.) The 3Com Layer 3 switching methodology is fundamentally unaffected by the choice between Gigabit Ethernet and ATM. In fact, the backbone could very well support both.

Phase 3

The third phase (Figure 7) completes the evolution of the backbone, and introduces policy services into the infrastructure. Such policy, administered by the network manager, may extend as far as the desktop, enabling network

access and signaling mechanisms for CoS and QoS. With the infrastructure distributed, yet overlaid with a logical management structure, performance metrics can be tuned and modified, giving greater viability to Service Level Agreements (SLAs). The legacy backbone has now been entirely eliminated and replaced by higher-performance Gigabit Ethernet, ATM, or both. The new backbone is inherently scalable and is ready for any future network evolution. Though beyond the scope of this paper, the core network will also become the termination point for virtual private networks (VPNs) as remote offices access the corporate infrastructure via the Internet. For more information on VPNs, see the white paper "Private Use of Public Networks," available at www.3com.com.

Future Trends

The Layer 3 switching solution does not stop here. Expect more Layer 4 capabilities to become available, handling advancements in

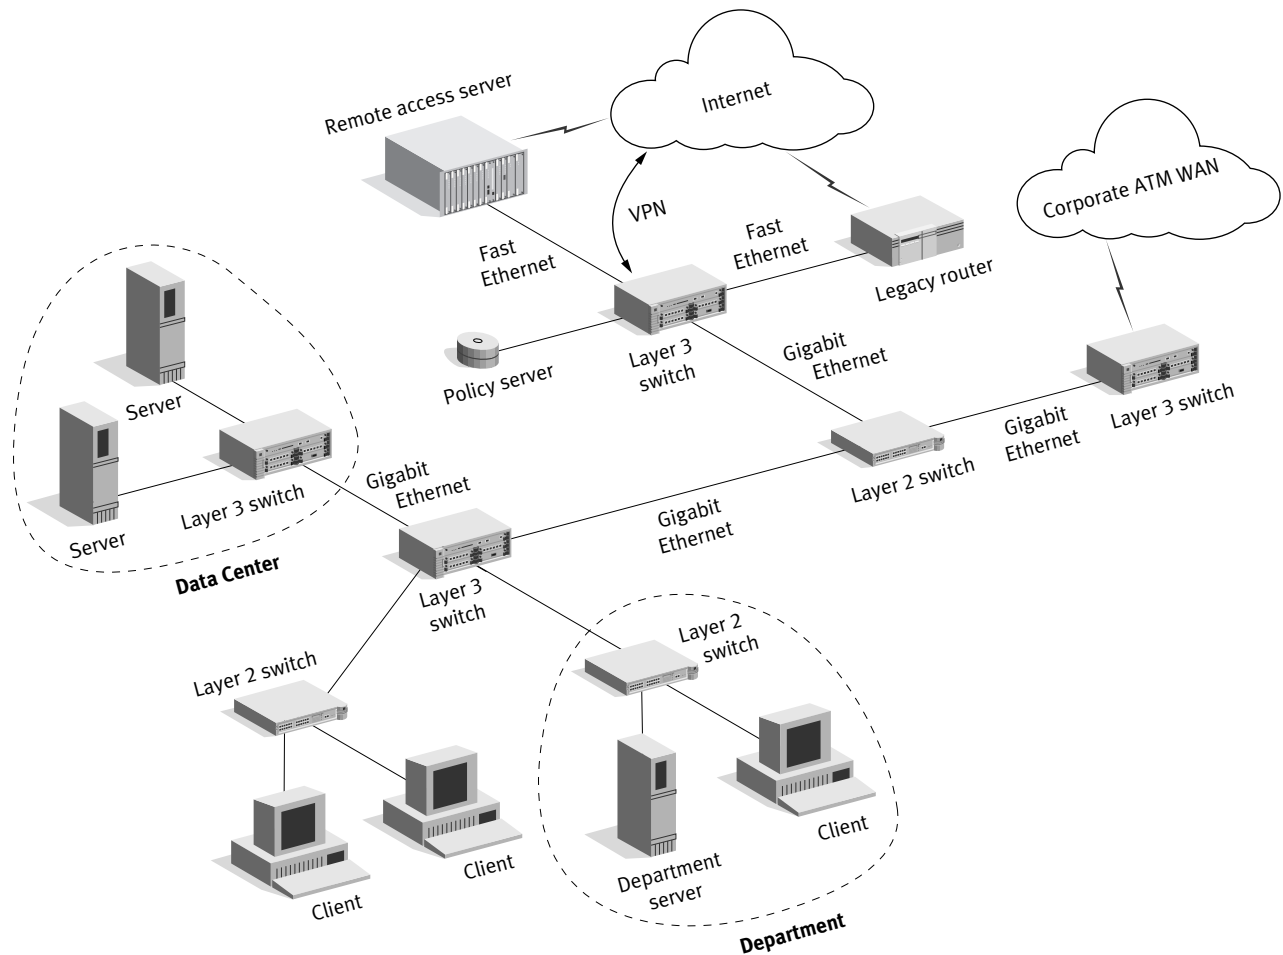


Figure 7. *Backbone Upgrade Complete; Policy Enabled*

middleware, along with providing more efficient Web server load balancing and caching. Directory-enabled networks will radically simplify the management paradigm using Layer 3 switching as a key delivery mechanism. VPNs will become more tightly coupled with the enterprise, interfacing more closely to the Layer 3 infrastructure. VPNs will have an increasingly significant role within the corporate intranet, requiring more security capabilities in the Layer 3 switch.

Conclusion

With a bit of knowledge, the Layer 3 jungle doesn't look so bad after all. In fact, Layer 3

switching is the natural evolution of networking technology and an enabling platform for next-generation applications. This progression represents the erosion of networking complexity, backed by increasing performance and decreasing cost. A Layer 3 switch turns out to be a well-known technology, not some entirely new model. But let the buyer beware. What looks like a true Layer 3 switch may not be one at all, so it is safest to invest in a product that was born as a true Layer 3 switch. With the advent of Layer 3 switching, the network is no longer a "Temple of Doom." Instead, it can fulfill its promise as a key element of enterprise business success. ▣



3Com Corporation

P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, CA
95052-8145
Phone: 1 800 NET 3Com
or 1 408 764 5000
Fax: 1 408 764 5001
World Wide Web:
<http://www.3com.com>

Asia Pacific Rim

Sydney, Australia
Phone: 61 2 9937 5000
Fax: 61 2 9956 6247
Melbourne, Australia
Phone: 61 3 9866 8022
Fax: 61 3 9866 8219

Beijing, China
Phone: 86 10 68492 568
Fax: 86 10 68492 789

Shanghai, China
Phone: 86 21 6350 1581
Fax: 86 21 6350 1531

Hong Kong
Phone: 852 2501 1111
Fax: 852 2537 1149

India
Phone: 91 11 644 3974
Fax: 91 11 623 3192

Indonesia
Phone: 62 21 572 2088
Fax: 62 21 572 2089

Osaka, Japan
Phone: 81 6 536 3303
Fax: 81 6 536 3304

Tokyo, Japan
Phone: 81 3 3345 7251
Fax: 81 3 3345 7261

Korea
Phone: 82 2 3455 6300
Fax: 82 2 319 4710

Malaysia
Phone: 60 3 715 1333
Fax: 60 3 715 2333

New Zealand
Phone: 64 9 366 9138
Fax: 64 9 366 9139

Philippines
Phone: 632 892 4476
Fax: 632 811 5493

Singapore
Phone: 65 538 9368
Fax: 65 538 9369

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

Thailand
Phone: 662 231 8151 5
Fax: 662 231 8158

3Com Benelux B.V.

Belgium
Phone: 32 2 725 0202
Fax: 32 2 720 1211

Netherlands
Phone: 31 346 58 62 11
Fax: 31 346 58 62 22

3Com Canada

Calgary
Phone: 1 403 265 3266
Fax: 1 403 265 3268

Edmonton
Phone: 1 403 423 3266
Fax: 1 403 423 2368

Montreal
Phone: 1 514 683 3266
Fax: 1 514 683 5122

Ottawa
Phone: 1 613 566 7055
Fax: 1 613 233 9527

Toronto
Phone: 1 416 498 3266
Fax: 1 416 498 1262

Vancouver
Phone: 1 604 434 3266
Fax: 1 604 434 3264

3Com Eastern Europe/CIS

Czech/Slovak Republics
Phone: 420 2 21845 800
Fax: 420 2 21845 811

Hungary
Phone: 36 1 250 8341
Fax: 36 1 250 8347

Poland
Phone: 48 22 6451351
Fax: 48 22 6451352

Russia
Phone: 7 095 258 09 40
Fax: 7 095 258 09 41

3Com France

Phone: 33 1 69 86 68 00
Fax: 33 1 69 07 11 54

3Com GmbH

3Com Austria
Phone: 43 1 580 17 0
Fax: 43 1 580 17 20

Berlin, Germany
Phone: 49 30 3498790
Fax: 49 30 34987999

Munich, Germany
Phone: 49 89 627320
Fax: 49 89 62732233

Switzerland
Phone: 41 844 833 933
Fax: 41 844 833 934

3Com Iberia

Portugal
Phone: 351 1 3404505
Fax: 351 1 3404575

Spain
Phone: 34 1 509 69 00
Fax: 34 1 307 79 82

3Com Ireland

Phone: 353 1 820 7077
Fax: 353 1 820 7101

3Com Latin America

U.S. Headquarters
Phone: 1 408 326 2093
Fax: 1 408 764 5730

Miami, Florida
Phone: 1 305 261 3266
Fax: 1 305 261 4901

Argentina
Phone: 54 1 312 3266
Fax: 54 1 314 3329

Brazil
Phone: 55 11 246 5001
Fax: 55 11 246 3444

*Chile (also serving
Bolivia and Peru)*
Phone: 56 2 633 9242
Fax: 56 2 633 8935

Colombia
Phone: 57 1 629 4847
Fax: 57 1 629 4503

Mexico

Phone:
52 5 520 7841/7847
Fax: 52 5 520 7837

Peru

Phone: 51 1 221 5399
Fax: 51 1 221 5499

Venezuela

Phone: 58 2 953 8122
Fax: 58 2 953 9686

3Com Mediterraneo

Milan, Italy
Phone: 39 2 253011
Fax: 39 2 27304244

Rome, Italy
Phone: 39 6 5279941
Fax: 39 6 52799423

3Com Middle East

Phone: 971 4 319533
Fax: 971 4 316766

3Com Nordic AB

Denmark
Phone: 45 48 10 50 00
Fax: 45 48 10 50 50

Finland
Phone: 358 9 435 420 67
Fax: 358 9 455 51 66

Norway
Phone: 47 22 58 47 00
Fax: 47 22 58 47 01

Sweden
Phone: 46 8 632 56 00
Fax: 46 8 587 057 99

3Com Southern Africa

Phone: 27 11 807 4397
Fax: 27 11 803 7405

3Com UK Ltd.

Edinburgh
Phone: 44 131 240 2900
Fax: 44 131 240 2903

Manchester
Phone: 44 161 873 7717
Fax: 44 161 873 8053

Marlow
Phone: 44 1628 897000
Fax: 44 1628 897003

To learn more about 3Com products and services, visit our World Wide Web site at <http://www.3com.com>.

© 1998 3Com Corporation. All rights reserved. 3Com and LANplex are registered trademarks and CoreBuilder is a trademark of 3Com Corporation or its subsidiaries. AppleTalk is a trademark of Apple Computer. VINES is a trademark of Banyan Systems. DECnet is a trademark of Digital Equipment Corp. IPX and NetWare are trademarks of Novell. XNS is a trademark of Xerox. Other brand or product names may be trademarks or registered trademarks of their respective owners. All specifications are subject to change without notice.