The Technology Guide Series
techguide.com

Layer 3 Switching

Re-Inventing the Router

This guide has been sponsored by

Cisco Systems

Intelligent Networking

Your enterprise network is in a state of perpetual growth and evolution. Increasing numbers of users are placing demands on your infrastructure, and traffic volume is on the rise. You have to accommodate new applications that your business depends on, such as video to the desktop, voice communications, and multimedia, which squeeze bandwidth and network resources. At the same time, you must structure and prioritize the flow of data to meet specific user needs, deliver to quality of service expectations, and sustain business-critical security standards. So, how can you ensure that your network is intelligent and scalable enough to consistently meet your growing needs?

Intelligent Networking from Cisco

Cisco is committed to helping enterprises build intelligent network infrastructures based on new technologies that empower your applications, users, and business. Cisco's approach is grounded in an in-depth understanding of the challenges that network managers face and three crucial building blocks that form the foundation of a truly intelligent network:

■ ***Intelligent devices*** such as campus switch routers and multilayer switches classify data in the network infrastructure and enable you to prioritize business-critical application data. And, with policy services and policy management under CiscoAssure Policy Networking, you gain the control you need to enforce network policy and maintain a more efficient network that runs on "smart bandwidth," not just more bandwidth.

■ ***Scalable performance*** enables enterprise networks to adapt to changing business requirements and remain cost-effective. Cisco's industry-proven solutions ensure maximum scalability, empowering you to manage the growth and performance of your network.

■ ***Network services*** embedded in Cisco IOS software® assist in delivering complete solutions to many network problems and performance issues. These services include Quality of Service (QoS), security, multiprotocol routing, multicast, reliability, and management. A Cisco intelligent network incorporates these network services into each layer, providing you with the full benefits of end-to-end service.

Comprehensive Enterprise Network Solutions

Cisco's intelligent networking solutions incorporate best-of-class products that meet your campus core and backbone needs by delivering the highest levels of performance and scalability. Now you can be assured that your business-critical applications are supported by the most intelligent networking devices available.

Cisco Systems

Visit ATG's Web Site

to read, download, and print

all the Technology Guides

in this series.

www.techguide.com

"The significant problems we face cannot be solved by the same level of thinking that created them."

Albert Einstein

Table of Contents

About the Editor...

Jerry Ryan is the Vice President of Editorial Development for the Technology Guides on Communications and Networking. Mr. Ryan is also a principal at ATG. Mr. Ryan has developed and taught many courses in network analysis and design for carriers, government agencies, and private industry. He has provided consulting support in the area of WAN and LAN network design, negotiation with carriers for contract pricing and services, technology acquisition, customized software development for network administration, billing and auditing of telecommunication expenses, project management, and RFP generation. He was the president and founder of Connections Telecommunications, Inc., a Massachusetts-based company specializing in consulting, education and software tools which address network design and billing issues. Mr. Ryan is a member of the Networld+Interop Program Committee. He holds a B.S. degree in electrical engineering.

The first generation of internetworks that were built in the early- to mid-1980s are now nearing the end of their life cycle. Today we find that shared LANs are not fast enough, servers have been consolidated into data centers, WAN speeds are increasing, voice and data networks are converging, and scalability has become an issue with the expanding PC population. Network availability is a critical success factor for mission-critical business systems, especially when they combine public and private networks. To respond to these many pressures, new technologies, each with its own terminology, have been introduced for high-speed LAN interconnection. This is resulting in serious confusion among potential buyers. Terms such as Layer 3 switching, multilayer switching, routing switches, switching routers, and gigabit switches regularly appear in the trade press, implying both a lack of standards and the inevitability of severe compatibility problems. While it is clear that a new generation of internetworking is now available and that users will benefit from innovative products and powerful new services, getting back to basics and de-mystifying the hype has never been more necessary. This Technology Guide describes the key functions and characteristics of a switched internetwork, clarifies what layer 3 switching is and is not, and demonstrates why routing will continue to be a part of future networks. The relationships among switching, route processing, management, and intelligent network services will also be analyzed.

Introduction

It was not so long ago that switches were used in telecommunications primarily for placing telephone calls. Dialing a telephone number activated a series of switches to set up a voice path that could be as simple as the next office or as complex as a multinational conference call. Internetworks and the Internet are beginning to provide similar services for PC workstations, servers, and mainframes. The primary goal of any data network provider is to eliminate geographic and media constraints on connectivity while maintaining control over resources and costs. Embedding Layer 3 Switching into the network is being promoted as the best way to achieve this.

Network planners and buyers are being bombarded with new terms that include the word “switch,” many of which claim to be breakthrough technologies. Layer 2, 3, and 4 switching, multilayer switching, LAN switching, routing switches, and switching routers are typical examples. Making sense of these terms and what they really represent, identifying the pros and cons of each approach and establishing a common frame of reference for network design are important tasks that should be a part of any purchasing decision.

This Technology Guide examines the concepts and realities of Layer 3 switching and the requirements common to any modern internetwork are identified. Layer 3 switches are compared to traditional multiprotocol routers. It is demonstrated that Layer 3 Switching is simply a re-invention of the router using new switch-based technologies. This Guide also reviews the basic data forwarding, route processing, and value-added functions that are required of any intelligent network node.

Next Generation Internetwork Requirements

Figure 1 illustrates, in simplified form, the major elements of any internetwork: hosts (workstations, servers, etc.), network nodes (hubs, bridges, routers, and switches), communication links, and management systems. Together, these form an enterprise infrastructure that can be viewed as multiple logical networks (i.e., the specific networks as seen by the users), each of which may have different functions, features, and qualities.

Networks are often organized along corporate geographic lines but may also be partitioned by user group, application type, or security policy. Complications arise when configurations change, when inter-enterprise connections are implemented, or when different technologies must co-exist. Network nodes clearly have to work together, so it is important that Layer 3 Switches be able to perform the same tasks as traditional routers.

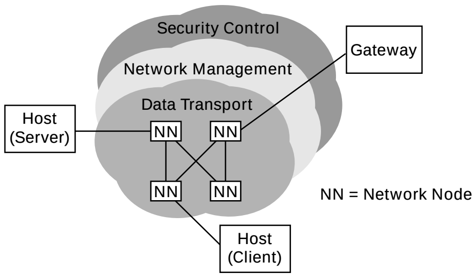

Figure 1: Basic Elements of an Internetwork

TCP/IP-based networks are now beginning to replace IBM's SNA as the corporate "nervous system." These include Intranets for internal use, connections to the worldwide public Internet for customer and public access and "extranets" for secure remote access via the Internet. Since these TCP/IP networks are mission critical, Layer 3 Switching must support a variety of configurations and must include all the manageability, reliability, and serviceability features that are assumed of routers.

High-speed switching, although a very important ingredient of any next generation internetwork, is far from being the only functionality required. A full-function network node must also include console support, application-awareness, access and usage controls, accounting features, and element management support. These functional requirements can be divided into the following five categories.

1. Network Integration and Interoperability: To be able to co-exist in an internetwork, Layer 3 Switches must provide functionality equivalent to today's routers and allow integration of network resources into a single physical network. Support for multi-user integration (independent user groups) and multiservice integration (more than one type of data) are also desirable. Layer 3 Switching can be implemented by updating traditional routers with new forwarding technology, thereby avoiding the need for interoperability testing, etc.

2. Capacity and Scalability: Any internetwork can be engineered either by over-provisioning the network (guaranteeing excess bandwidth) or by controlling access to resources on the basis of a pre-defined service allocation policy. Increasing the speed of the network (such as moving an Ethernet LAN from a traditional shared network to Fast Ethernet and soon to Gigabit Ethernet) demands that equivalent increases in switching speeds be provided. Layer 3 Switches support packet-forwarding throughputs in the millions of packets per second (as compared to a million for traditional general-purpose routers). Aggregate performance

is one of the only real differences between the new Layer 3 Switching techniques and the traditional software-driven routing mechanisms.

3. Quality of Service: Every network exhibits specific QoS characteristics that result from the physical media being used, the design of the network, and the network node equipment. Recently, however, a need to explicitly define a “contract” for the quality of service has emerged. Controlling latency and jitter is important in voice networks, for example, especially if the network is shared by different applications. Layer 3 Switching increases the packet forwarding speed of the node, thereby minimizing congestion.
4. Control and Management: Network management includes both element and system management and also covers security controls, quality of service control, name and address administration, policy controls, etc. Any network benefits from a well-developed network management system that facilitates dynamic operation of the network, supports automatic fault avoidance, applies policies when available, and controls network operations to the best advantage of all its users. Whether control and management tasks are fully distributed, localized, or are centralized is of less concern than are the improvements in overall effectiveness. Operation, administration, maintenance, and security of the internetwork is much less effective when protocols and mechanisms are incompatible or are not consistently implemented. Layer 3 Switches must support the same network management services and interfaces as any other network component.
5. Cost and Economies of Scale: The cost of a network is always of concern to potential buyers. Although the direct cost of network equipment is the most visible component, it is the total cost of ownership that should be the most important criteria. Designing the flexibility to adapt to changing circumstances while also avoiding disruptive network upgrades is one area where savings can also be achieved. Users move frequently, servers are being re-centralized, traffic patterns are changing (from 80% local/20% remote to the opposite) and data content is more complex and time-sensitive, all of which simply mean that change will be continuous. Layer 3 Switches and routers must be equivalent in this area.

In short, the key to the next generation of internetworks is the re-development of the functionality of Layers 2 and 3 to take advantage of the new high-speed switching hardware, the advances in protocol software, and the intelligent network processing features that are considered to be standard in traditional routers.

The Layering of Network Functions

A brief review of the key OSI layers, and of the Layer 3 in particular, can help to put Layer 3 Switching into context and illustrate the role of switching in network processing. Layering, which was first used to divide network functions into independent subsets, is best described by the OSI Reference Model. Table 1 on the following page illustrates the elements of switching that are present in any network protocol stack.

Table 1: Overview of Internetworking Functionality

Layer	Purpose of the Layer	Role of Switching
(5-7) Application, Presentation, and Session	Defines user-oriented services such as file transfer, messaging and transaction processing; provides for structuring applications, coding the data, and exchanging information	Application switching (e.g., e-mail forwarding); gateways between different application types; support for management functions; selection of destinations for messages
(4) Transport	Delivery of data to applications, division of messages into packets	Directs the messages to the specific destination application or protocol type
(3) Network	End-to-end communications through one or more sub-networks; selects optimal routes; controls loops; manages addressing	Forwards packets through an interconnected set of networks
(2) Data Link	Transfer of frames across a single network link such as a LAN; manages contention	Controls switched circuits, switched LANs, and recovers from link errors
(1) Physical	Transmission over a physical circuit including physical connectors, bit encoding, etc.	Circuit switching as is used for telephony and port switching for LAN physical media

As can be seen from the table, each layer plays a specific role in the process of communicating and each provides mechanisms that contribute to the ultimate goal of application interworking. Layer 3 is especially important, however, because it overcomes the technical differences among networks (between LANs and WANs, for example) and provides physical connectivity from end-to-end. This relieves Layer 4 (and the applications) from having to know about the physical configuration of the network and the complexities of each of its components.

Switching can be defined as the process of forwarding units of data towards its destination (in the form of cells, frames, or packets). Each switch in the network must make forwarding decisions based on the information that is available to it (with a higher layer having more information than that of the layers below). Knowing at least one data link that is headed in the right direction is the minimum; understanding the

traffic expectations, service level requirements, and security restrictions would allow a much more informed switching decision.

Devices that perform switching-related functions are usually classified by the layer in which they operate. Figure 2 illustrates the relative positioning of the traditional categories of network node. Complications in terminology arise when a single device combines more than one category or uses information from multiple layers.

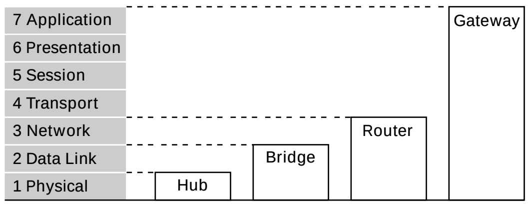

Figure 2: Types of Internetwork Nodes

Figure 2 illustrates the positioning of hubs, bridges, and routers in the protocol hierarchy. The forwarding functions of each type of node are introduced in the following sections.

- Layer 1 Nodes (Hubs / Repeaters): Any device that simply copies bits that arrive on an interface to all other available interfaces is performing the role of a repeater. Repeaters regenerate the electrical and timing characteristics of signals and can convert between media types (such as copper wire to optical fiber). Hubs are multiport repeaters that can operate as multiple two-port repeaters or as one or more multi-drop LANs. The inputs and outputs of a hub can be “wired up” permanently or can be configured through software port assignments (i.e., an electronic patch panel). Users connected to a hub group are all on the same LAN and must contend with each other for bandwidth. Broadcast traffic is not filtered by a hub.

- Layer 2 Switches (Multiport Bridges): Bridges offer a frame forwarding service based on the physical addresses that are available as part of Layer 2 (i.e., the MAC address of the destination) as well as performing the signal regeneration functions of a repeater. A bridge monitors the traffic to learn which addresses exist on which ports and then builds a table of forwarding rules to control the switching process. Bridges must also identify and eliminate potential data loops (using the spanning tree algorithm). A Layer 2 Switch functions as a multiport bridge. An internetwork built entirely out of Layer 2 Switches appears as a single large network with a "flat" address space. Layer 2 Switched networks have limited flexibility and scalability.
- Layer 3 Switches (Routers): Routers offer a packet forwarding service based on the logical destination addresses that are available as part of Layer 3 (as well as providing bridging and repeating functions). Routers cooperate (using routing protocols) to discover all possible paths between end systems and to maintain tables of address to port mappings. Routers view the network as a hierarchically ordered set of address blocks, not as a huge set of randomly assigned addresses. Routers use Layer 3 header information to make selective forwarding decisions, allowing broadcast and multicast traffic to be suppressed. Multiple active paths between communicating hosts can be supported. Conversions between different types and speeds of data link (such as LAN-WAN conversions) are supported by routers. Routers are also intelligent enough to provide functions such as access controls based on the type of application protocol being carried and various forms of packet filtering.
- Layer 7 Nodes (Applications and Gateways): Applications such as electronic mail can include switching functions in the Application Layer (messages are forwarded through post offices). An Application Gateway supports interworking among dissimilar application environments, and is mostly concerned with the transformation of business functions to operate over different protocols. Gateways are usually application-specific and must be custom designed. The need for gateway capabilities can also be demonstrated by electronic mail systems when different protocol suites (X.400 and SMTP, for example) need to exchange mail. Other services that support the operation of switching, such as system management and directories, are also positioned at the Application Layer. Other forms of switching (application and presentation context switching, for example) are defined in the OSI Reference Model but are not widely implemented.

One of the goals of every layered architecture is the separation of the layers: a lower layer should not need to know about the layers above it in order to perform its tasks. For example, a Layer 2 Switch should not require TCP port numbers or network addresses in order to forward frames on a data link. The higher you are in the hierarchy the more information will be available to work with and the more control you will be able to exert. For example, Layer 2 performs its tasks using knowledge derived from a frame header that includes link addresses, while Layer 4 understands that one of its segments is carrying a file (a long flow), another is carrying a short e-mail message, and a third could be real-time voice (a time-sensitive stream of data). Thus, "value added services" are possible that would not be available if only Layer 2 switching is used. This is the same reason why bridged networks were originally augmented with traditional routers.

Network technologies have improved dramatically over the last decade. More powerful features are available, higher quality transmission techniques are being used, and costs are significantly lower than ever before. The introduction of packet switching in the 1970s and the popularization of LANs in the 1980s led to the concept of a "network of networks", or in current terminology, an internetwork. The rapid growth of the Internet in the early 1990s established TCP/IP as the standard protocol "stack" for data transfer and, more recently, for the convergence of voice, data, and video systems. More recently, Ethernet LANs operating at a speed of one gigabit per second have become feasible and this has created a need for equivalent increases in packet forwarding rates. This is being achieved by re-inventing the router as high-speed Layer 3 Switches.

Layer 3 Switching in an Internetwork

Layer 3 Switching has become a popular term in the technical press, with the implication being that it represents a breakthrough technology that is destined to replace conventional routers. The reality is that a Layer 3 Switch is simply a class of high performance router that is optimized for the campus LAN or intranet. Aggregate performance is the primary difference between a conventional router and a Layer 3 Switch.

Switches and routers can be compared by examining each of the major functional components. Figure 3 illustrates the switching (packet forwarding) function, the processing of routing information, the provision of intelligent network services, and node management.

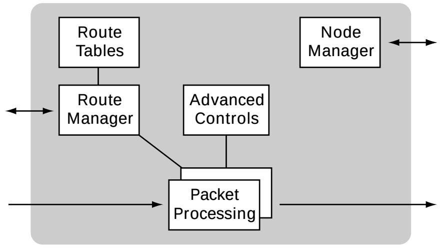

Figure 3: Layer 3 Node Components

Packet Processing

The task of a switching "engine" (the actual switch plus the decision processor) is to examine incoming packets, extract the destination address(es) and other control information, compare the address to the contents of a routing table, re-write the packet control data, and then submit the packet to the appropriate output interface for transmission. Switching performance is an indication of how fast it can dispatch packets towards their destinations. The measure is usually the aggregate number of packets that the device can switch in and out over a fixed period of time.

The faster the packets arrive the faster they need to be processed. Otherwise, traffic congestion occurs, performance degrades, and ultimately the network service fails. The requirement is to minimize the time it takes to complete all the processing that is needed with the goal being forwarding at "wirespeed" (the data is forwarded as fast as it arrives). A guarantee that switching delays will be consistent from one packet to the next is also desirable for time-sensitive data such as voice or video.

A number of techniques for increasing Layer 3 performance are being used:

- Increasing the processing speed by using specialized high-performance hardware;

- Reducing processing overhead by eliminating certain capabilities or avoiding features of unusually high complexity (supporting only IP protocols, for example);
- Avoiding the need for full packet analysis and processing at every node in the path (such as by using combinations of Layer 2 and 3 switching); and
- Taking advantage of any external knowledge that may be available (such as identification of flows, prioritization of processing, etc.).

Only the first of these provides a transition that is completely compatible with exiting equipment and is fully transparent to the end user.

If every packet is examined completely and forwarded as a separate, independent unit of data, the device is performing "packet-by-packet" processing. Each packet is individually processed as it is received and then forwarded whenever a suitable outgoing link is available. Processing for features such as access controls and filtering would be performed on each packet. Packet-by-packet switching is relatively straightforward, fully supports value-added services, is well-understood, and has been well-tested in current products.

Packet-by-packet processing and Layer 3 Switching are actually equivalent, with the only major difference being the physical implementation. In conventional routers, processing takes place in microprocessor-based engines whereas a Layer 3 Switch uses ASIC (Application Specific Integrated Circuit) hardware. Layer 3 Switches can achieve forwarding rates that are much higher than is feasible with a central microprocessor (millions of packets per second can be achieved).

A second approach to switching is called "cut-through" processing. Cut-through processing takes advantage of added knowledge to reduce packet processing overheads. The first packet is fully processed at Layer 3 with subsequent packets either forwarded at Layer 2 or handled using tag switching mechanisms. The assumption is that the processing overhead of Layer 3 can be reduced if multiple packets can be associated with a specific data flow (defined as having the same source and destination addresses, protocol port number, type of service, and input interface port) or a label of some kind that associates the packets with each other.

In the long run, improvements to the raw speed of Layer 3 forwarding needs to be balanced against the provision of features that depend on Layer 3 information. The combination of wirespeed forwarding using ASIC processors and the addition of packet labels may prove to be satisfactory for most emerging applications.

Route Processing and Management

Every Layer 3 node (either a conventional router or a Layer 3 Switch) must be able to establish and continuously update the list of available routes to any given destination. This information is fundamental to the switching decision process (that is, selecting the output interface for the packet). Routing tables may be initialized before traffic flow begins but must be kept synchronized with changes to the network topology as they occur. Dynamic management of routing information has traditionally been the task of routers, regardless of the mechanisms used by the switching engine, and is a software-based activity.

Changes in the topology of a network can be caused by link failures, device failures, additions and deletions to the network, and changes in quality of service metrics. In a large network, the number of routes and the number of route-affecting events can be very large, with changes needing to be propagated very rapidly. The Internet is the prime example of a network that can stress routing protocols to their limits, but this is also increasingly true for high-speed local networks. Updates to routing tables, if not completed quickly, seriously impact the overall performance of an internetwork.

A number of packet routing protocols have been developed that may need to be supported by a Layer 3 node. These include RIP, the first dynamic routing protocol to be deployed, as well as RIP II, OSPF, IGRP and EIGRP—all of which could be deployed in an IP-based network. Routers also support load sharing across multiple parallel links to increase total capacity and provide fault tolerance. More recently, extensions to these protocols have been introduced for multicast transport services (with Protocol Independent Multicast, or PIM, emerging as the protocol of choice). Routing protocol support is a basic requirement for any Layer 3 Switch. The stability, quality, and maturity of the implementation is also an important buying criteria.

In most real-world networks, more than one routing protocol will be in use. In this “multilingual” environment, different route processing subsystems must cooperate to ensure global optimization or resources and to minimize conflicts (after all, they are all accessing the same inter-node links). The ability to jointly optimize across different protocol suites can allow the network designer to tailor their network structures and operational characteristics to suit individual needs.

A Layer 3 Switch cannot avoid the need for routing information. Although packet switching performance may grab the headlines and allow easy comparisons, route information processing and the management of distributed routing tables are equally essential to internetwork operation.

Network Operations Support

The ability to forward packets quickly to the right destination over optimal routes is only part of the overall internetworking requirement. Value-added features that are often taken for granted in today’s routers would need to be re-created for each new internetwork technology. This includes QoS control, mobility, software support and distribution, reliability, and addressing support.

While intelligent network services permit applications to run more efficiently on the network, they are also typically software-dependent and processing-intensive. A few examples of this type of service include:

- Frame translation: Routers provide for conversions from one type of frame to another (such as Ethernet to FDDI, etc.) and for different link speeds;
- Advanced queuing: A variety of queuing mechanisms are often available to assist with traffic prioritization and other QoS functions;
- Load balancing: Multiple links between routers can be used to provide redundancy and load sharing features;
- Hot standby routers: Traffic can be transparently routed around a failed Layer 3 node with minimal application downtime and no impact on the end-user station through the use of a protocol allowing the routers to cooperate; and
- Support for mobility: This includes Local Area Mobility and Dynamic Host Configuration Protocol.

Since virtually all of these features are used regularly in today’s networks, they must remain intact (and use the same standards) during and after the transition to Layer 3 Switching.

Management and Security Services

Conventional routers routinely add value through their support of network management and security services, both locally at the network node and system-wide. Manageability enhancements can reduce the total cost of ownership, support the network operators goal of “zero downtime,” and provide accounting and performance data. Security, on the other hand, is a defensive measure that helps to eliminate inappropriate use of resources, intentional disruption of services, and theft of information.

To detect or solve a network problem, it is often necessary to access a network node from a remote location to gather diagnostic information or to make configuration changes. Access to routers via the Telnet protocol is usually available. Multiple inbound and outbound Telnet sessions are often used to reach a device that needs to be modified. In making configuration changes, it is imperative that the change take effect immediately on the node in question, without having to reboot or partially reboot the device (which would be service affecting).

Support for other TCP/IP-based protocols can also be invaluable. The Simple Network Management Protocol (SNMP) and File Transfer Program (FTP) are especially important. SNMP is needed to collect management information such as errors and events for a central manager and has become a standard for most networks. FTP is used in conjunction with the distribution of software or configuration files, making central administration of the network much easier.

Other intelligent network services that can improve the total cost of ownership for a network include:

- simplifying the processes of moves, adds, and changes for local LAN segments, for network-based resources, and for end users;
- tools for local and/or remote troubleshooting and debugging;
- logging of warnings, errors, and debug information for collection and off-line analysis; and
- network time-of-day synchronization.

Security support embedded in the network nodes can be used to protect the network itself, to prevent unauthorized access to information and systems, and to ensure privacy. Examples of mechanisms used in routers that would need to be included in Layer 3 Switches are:

- Access control lists: Can be used to filter or deny traffic based on tests of the packet (such as protocol type, addresses, special processing);
- Host access protection: Provided using security mechanism such as encrypted passwords and configuration files, authentication of user logon attempts, etc.; and
- Node protection: Provided by using console passwords, privileged mode operation, etc. to discourage hackers.

The essence of the argument is that the software in each router node combines to form a distributed “internetwork operating system” that cannot be easily replicated in a new product. In order to function optimally, the network must be compatible across all its elements. Thus, the decision of high-speed switching versus traditional software-based routing is of less concern than making sure all the advanced features are available, interoperable, and thoroughly tested.

None of the intelligent network services are unusually difficult to implement in a new type of node. They do, however, involve processing support at the node and must be integrated with the other functions.

Extending existing routers to serve as Layer 3 Switches (i.e., re-inventing today's routers) allows all of the existing investments and development experiences to be preserved, thereby reducing costs and allowing greater assurance of success.

The Benefits of Layer 3 Switching

Achieving an acceptable balance between the packet forwarding capabilities of network nodes, the speeds achievable on the data links (local and wide area), the resiliency and fault tolerance of the configured network, the advanced processing features embedded in the internetwork, and the total cost of ownership is always a major challenge for planners and designers. This challenge is even greater during periods of change when new applications are being introduced to older networks or when combinations of old and new technologies are being used.

In this Technology Guide we have defined Layer 3 Switching as being the name for a network node that can meet the processing requirements of high-speed networks, with or without the addition of mechanisms such as recognition of traffic flows, substitution of labels for destination addresses as the basis for forwarding decisions. In this sense, Layer 3 Switching is simply a new name for the next generation of LAN routers. The defining characteristics of these devices is their ability to work with switched LANs, to support Gigabit Ethernet speeds, and to handle more complex traffic patterns efficiently. The benefits of Layer 3 Switching, however, go beyond simply higher speeds.

Rapid growth in the deployment of the Internet and intranets in general, and subsequently in their usage, has caused a major shift in the way people use computing both at the office and at home. There have been massive increases in demand for network bandwidth, performance, predictable quality of service, access and traffic controls, and support for multimedia content. Layer 3 Switching can provide the basis for intelligent support for this new generation of multiservice, multifunction networks.

Layer 3 Switching can offer several significant benefits for the internetwork service provider, including the following six items.

- Wirespeed packet forwarding

  The introduction of higher speed LANs into the network infrastructure forces switching nodes to work harder. The key is not to re-invent the router using new technologies simply because higher performance is needed. An approach that adopts conventional routers as the basis for increasing the forwarding speeds, for extending and improving the route decision processes, and for developing more intelligent network services offers a lower risk evolution strategy.

  Packet forwarding at the speed of the data links is an essential requirement for next generation networks, but it is not essential to invent new ways of doing routing in order to achieve this. Being able to use both conventional routers and Layer 3 Switches in the same network avoids the requirement for expensive "fork-lift" upgrades and allows a gradual transition based on the performance requirements of each user.

- Standards-based technology

  The adoption of gigabit speed LANs will neither be an overnight event nor will such speeds be needed for every network host. No sudden, widespread need for multimedia services or advanced traffic control is expected in the near future. Nevertheless, compatibility and interoperability become more critical whenever a network includes multiple generations of equipment or uses products from different vendors. Development of Layer 3 Switches based on industry standard protocols will avoid most issues related to compatibility and investment protection.

Not all the different approaches to Layer 3 Switching are standards-based. Retention of core routing functionality combined with faster processing and the addition of new features is expected to be an approach that meets most user expectations.

- Scalability

  As has been stated previously, scalability is a major driving force for the introduction of high performance Layer 3 Switching. The expectation of continued growth in the number of internetwork users (especially on the Internet), the amount of complex traffic that will have to be carried, the global distribution of access points, and the large number of service providers to be interconnected all lead to one conclusion: an order of magnitude improvement in network switching will be needed.

  Layer 3 Switching takes advantage of the latest in hardware and software to provide the levels of scalability that are clearly going to be required.

- Control and management

  Layer 3 Switching requires only incremental additions to existing routers and LAN switches allowing all of the existing control and management services to remain intact.

- Support for “legacy” protocols

  It is well-recognized that TCP/IP has been adopted as the basic suite of protocols for internetworking and that local networks are converging on Ethernet standards. In most cases, the gains in processing efficiency that can be derived from reducing the “multi-” part of a multiprotocol router will be significant. It must be recognized, however, that this may not be as easy to accomplish as would be desirable and that support for legacy protocols in parts of the internetwork may continue to be needed.

  Layer 3 Switches that are derived from an existing router base can offer this type of support without requiring re-development of software for non-strategic protocols.

- Cost/performance gains

  The bottom line for most purchasers is the increased performance/cost ratio that can be achieved with Layer 3 Switching. This is primarily the result of adopting newer, higher performance hardware.

Summary and Conclusions

Layer 3 Switching is not a revolutionary new concept or even a break-through in networking technologies; rather, it refers to a new class of high performance routers that are optimized for the campus LAN or intranet. Today’s routers are simply being re-engineered to handle the high speeds of Gigabit Ethernet, to integrate data with voice and video, and to accommodate the increasingly distributed nature of the traffic flows.

Layer 3 Switches are re-invented routers that establish a new balance of power between the high-performance network hosts that are now being deployed, the high-speed LANs that are becoming available. They also exploit the latest advances in hardware, software, and transmission technologies to establish a new standard for cost/performance.

The techniques used with Layer 3 Switching offer definite improvements in the speed of packet forwarding but do nothing to reduce the requirement for route processing services, congestion avoidance mechanisms, access controls, and network element management. The combination of high performance packet switching, robust routing protocols, and intelligent network services, first introduced in traditional routers, continues to be the formula for internetworking success.

Glossary

Address—A means for mapping Network Layer resolution addresses onto media-specific addresses.

Aggregate Data—Data that is the result of applying a process to combine data elements. Data that is taken collectively or in summary form.

Aggregation—A token number assigned to an outside link by the border nodes at the ends of the outside link. The same number is associated with all uplinks and induced uplinks associated with the outside link. In the parent and all higher-level peer groups, all uplinks with the same aggregation token are aggregated.

Application Specific Integrated Circuit (ASIC)—A hardware technology where integrated circuits are manufactured for customized applications in order to reduce space. It is also manufactured for speed and security reasons.

Backbone—(1) The part of a network used as the primary path for transporting traffic between network segments. (2) A high-speed line or series of connections that forms a major pathway within a network.

Bridge—A device that connects and passes packets between two network segments. Bridges operate at Layer 2 of the OSI reference model (the data-link layer) and are insensitive to upper-layer protocols. A bridge will examine all frames arriving on its ports and will filter, forward, or flood a frame depending on the frame's Layer 2 destination address.

Cache—A high-speed dynamic memory used as a buffer between the CPU and physical disk storage to mitigate or eliminate potential speed differences between access times to physical disks and faster system memory. In storage arrays, cache implementation is usually non-volatile to ensure data integrity.

Central Processor Unit (CPU)—The computer.

Classification—The process of dividing a data set into mutually exclusive groups such that the members of each group are as “close” as possible to one another and different groups are as “far” as possible from one another, where distance is measured with respect to specific variable(s) you are trying to predict. For example, a typical classification problem is to divide a database of customers into groups that are as homogeneous as possible with respect to a creditworthiness variable with values “Good” and “Bad”.

Client/Server—A distributed system model of computing that brings computing power to the desktop, where users (“clients”) access resources from servers.

Configuration—The phase in which the LE client discovers the LE Service.

Connection—An ATM connection consists of concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

Cut-Through Switching—Refers to a method of Frame Switching where the switching device commences forwarding a frame after it has determined the destination port without waiting for the entire frame to have been received on the incoming port. Also known as on-the-fly switching.

Data Link—A layer of the OSI Reference Mode (Layer 2). Also used as the name for a connection between two network nodes.

Edge Device—A physical device which is capable of forwarding packets between legacy interworking interfaces (e.g., Ethernet, Token Ring, etc.) and ATM interfaces based on data-link and network layer information but which does not participate in the running of any network layer routing protocol. An Edge Device obtains forwarding descriptions using the route distribution protocol.

Electronic Commerce (EC)—The automated transaction of business, including the transfer of both information and funds, via computers.

End System—End-user device on a network. Also, a nonrouting host or node in an OSI network.

End User—Any customer of an interstate or foreign telecommunications service that is not a carrier, except that a carrier other than a telephone company shall be deemed to be an “end user” when such carrier uses a telecommunications service for administrative purposes. It is also a person or entity that offers telecommunications services exclusively as a reseller if all resale transmission offered by such reseller originate on the premises of such reseller.

Enterprise—A complete business consisting of functions, divisions, or other components used to accomplish specific objectives and defined goals.

Enterprise Network—(1) A geographically dispersed network under the auspices of one organization. (2) A complete business network consisting of functions, divisions, or other components used to accomplish specific objectives and defined goals.

Error Rate—The ratio of the number of data units in error to the total number of data units.

Ethernet—(1) A baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks operate at 10 Mbps using CSMA/CD to run over coaxial cable. Ethernet is similar to a series of standards produced by IEEE referred to as IEEE 802.3. (2) A very common method of networking computers in a local area network (LAN). Ethernet will handle about 10,000,000 bits per second and can be used with almost any kind of computer.

Ethernet LANs—An Ethernet LAN is a local networking technology that is based on IEEE 802.3 standards. It uses Carrier Sense Multiple Access (CSMA) with Collision Detection (CD) transmission techniques. Traditional Ethernet LANs operate at 10 Mbps over twisted copper wire, coaxial cable, and fiber optic cables. Fast Ethernet (100 Mbps) and Gigabit Ethernet (1 Gbps) are more recent versions of the basic Ethernet LAN.

Extranet—A collaborative network that uses Internet technology to link businesses with their suppliers, customers, or other businesses that share common goals.

Fault Tolerance—Generally, the ability to prevent a problem on a device from affecting other devices on the same port.

Forwarding—The process of sending a frame toward its ultimate destination by an internetworking device.

Frame—A logical grouping of information sent as a link-layer unit over a transmission medium. The terms packet, datagram, segment, and message are also used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.

Gateway—(1) A set of functions intended to facilitate electronic access by users to remote services and vice versa. Gateways are intended to provide a single source through which users can locate and gain access to a wide variety of service. Gateways typically offer a directory of services available through them, and provide billing for these services. (2) Technically, a Gateway is a hardware and/or software connection that translates information between two dissimilar protocols. The term has come to be used, however, to describe any mechanism for providing access from one system to another.

Gigabit Ethernet—A 1 Gbps standard for Ethernet.

Gigabits Per Second (Gbps)—Billion bits per second. A measure of transmission speed.

Intelligent Network Node—Any network element providing bridging, routing and/or switching capabilities that is also manageable and configurable.

Interface—(1) A connection between two systems or devices. In routing terminology, a network connection. Also, the boundary between adjacent layers of the OSI model. In telephony, a shared boundary defined by common physical interconnection characteristics, signal characteristics, and meanings of interchanged signals. (2) The point at which two systems or pieces of equipment are connected.

Internet Address—Also called an IP address. It is a 32-bit address assigned to hosts using TCP/IP. The address is written as four octets separated with periods (dotted decimal format) that are made up of a network section, an optional subnet section, and a host section.

Internet Protocol (IP)—A Layer 3 (network layer) protocol that contains addressing information and some control information that allows packets to be routed. Documented in RFC 791.

Internetworking—General term used to refer to the industry that has arisen around the problem of connecting networks together. The term can refer to products, procedures, and technologies.

Interoperability (IOP)— (1) The ability of equipment from different manufacturers (or different implementations) to operate together. (2) A term that implies that different vendor products of the same technology can successfully operate with each other.

Jitter—Analogue communication line distortion caused by a variation of signals from its reference timing positions. Jitter can also cause data loss, particularly at high speeds.

Layer—A level of the OSI Reference Model. Each layer performs certain tasks to move the information from sender to receiver. Protocols within the layers define the tasks for networks, but not how the software accomplishes the tasks. Interfaces pass information between the layers they connect.

Layer 3 Switching —The emerging Layer 3 switching technology integrates routing with switching to yield very high routing throughput rates in the millions-of-packets-per-second range. The movement to Layer 3 switching is designed to address the downsides of the current generation of layer 2 switches, which are functionally equivalent to bridges. These downsides for a large, flat network include being subject to broadcast storms, spanning tree loops, and address limitations that drove the injection of routers into bridged networks in the late 1980s. Currently, Layer 3 switching is represented by a number of approaches in the industry.

Link—Physical connection between two nodes in a network. It can consist of a data communication circuit or a direct channel (cable) connection.

Load Balancing—In routing, the ability of the router to distribute traffic over all its network ports that are the same distance from the destination address. It increases the use of network segments, which increase the effective network bandwidth.

Local Area Network (LAN)—(1) A privately owned data communications system that provides reliable high-speed, switched connections between devices in a single building, campus, or complex. (2) A network covering a relatively small geographic area (usually not larger than a floor or small building). Compared to WANs, LANs are usually characterized by relatively high data rates. (3) Network permitting transmission and communication between hardware devices.

Megabits Per Second (Mbps)—A digital transmission speed of millions of bits per second.

Multilayer Switching (MLS)—A network technology that combines Layer 2 and Layer 3 switching in a single network node.

Multiservice—An adjective describing a network or component that handles multiple types of data including voice, data, and video.

Network—A collection of computers and other devices that are able to communicate with each other over some network medium.

Network Address—(1) Also called a protocol address. A network layer address referring to a logical, rather than a physical, network device. (2) Numeric character string used to specify the location of the called customer.

Network Layer—Layer 3 of the OSI reference model. Layer 3 is the layer at which routing occurs.

Network Management System (NMS)—A system responsible for managing at least part of a network. NMSs communicate with agents to help keep track of network statistics and resources.

Network Node Interface (NNI)—Set of ATM Forum-developed specifications for the interface between two ATM nodes in the same network. Two variations are being developed; an interface between nodes in a public network (called a public NNI or just NNI) and an interface between nodes in a private network (called a private NNI or P-NNI).

Node—One of a number of UNIX Computers joined together to form a Staffware Network.

Node Management—The management of a network node.

OSI Reference Model—Seven-layer network architecture model of data communication protocols developed by ISO and CCITT. Each layer specifies particular network functions such as addressing, flow control, error control, encapsulation, and reliable message transfer.

Packet—(1) A logical grouping of information that includes a header and (usually) user data. (2) Continuous sequence of binary digits of information is switched through the network and an integral unit. Consists of up to 1024 bits (128 octets) of customer data plus additional transmission and error control information.

Packet Filtering—A second layer of filtering on top of the standard filtering provided by a traditional transparent bridge. Can improve network performance, provide additional security, or logically segment a network to support virtual workgroups.

Packet Forwarding—The function of moving data packets through a router from an input port to an output port. Packet forwarding is a basic function of a router.

Packet Switching—Type of data transfer that occupies a communication link only during the time of actual data transmission. Messages are split into packets and reassembled at the receiving end of the communication link. (2) A transmission technique that segments and routes information into discrete units. Packet switching allows for efficient sharing of network resources as packets from different sources can all be sent over the same channel in the same bitstream.

Points of Presence (POP)—A term used by Internet service providers to indicate the number of geographical locations from which they provide access to the Internet.

Policy—Organization-level rules governing acceptable use of computing resources, security practices, and operational procedures.

Private Network—A leased, private transmission network that links multiple locations of a company or other organization, using voice and/or data communications lines reserved exclusively for that company's traffic.

Protocol—(1) A formal description of a set of rules and conventions that govern how devices on a network exchange information. (2) Set of rules conducting interactions between two or more parties. These rules consist of syntax (header structure) semantics (actions and reactions that are supposed to occur) and timing (relative ordering and direction of states and events). (3) A formal set of rules.

Protocol Software—Any software that implements a protocol.

Protocol Stack—Related layers of protocol software that function together to implement a particular communications architecture. Examples include AppleTalk and DECnet.

Quality of Service (QoS)—Term for the set of parameters and their values which determine the performance of a given virtual circuit.

Repeater —(1) A device that regenerates and propagates electrical signals between two network segments. (2) Device that restores a degraded digital signal for continued transmission; also called a regenerator.

Route Processing—A processor board determining a router and running configuration, security, accounting, debugging, and network management processes. The CSC/3 and the CSC/4 are route processors that perform this.

Routing—The process of finding a path to the destination host. Routing is very complex in large networks because of the many potential intermediate destinations a packet might traverse before reaching its destination host.

Scalability—The ability to add capacity to the network element to accommodate growth.

Server—(1) A software application that responds with requested information or executes tasks on the behalf of a client application. Also, a network host, such as a web server, running a set of protocol server applications. (2) Any computer that allows other computers to connect to it. Most commonly, servers are dedicated machines. Most machines using UNIX are servers.

Signaling—(1) The process of sending a transmission signal over a physical medium for purposes of communication. (2) Method of communication between network components to provide control management and performance monitoring.

Simple Network Management Protocol (SNMP)—The Internet network management protocol. SNMP provides a means to monitor and set network configuration and runtime parameters.

Switch—In the context of Frame or LAN switching, this refers to a device which filters, forwards, and floods frames based on the frames destination address. The switch learns the addresses associated with each switch port and builds tables based on this information to be used for the switching decision. Some switches are high-speed implementations of bridges where switching decisions are made in silicon, usually an Application Specific Integrated Circuit (ASIC).

Switched LAN—Refers to a LAN implemented with frame switches.

Systems Network Architecture (SNA)—IBM's proprietary network architecture.

Transport Control Protocol/Internet Protocol (TCP/IP)—A protocol (set of rules) that provides reliable transmission of packet data over networks.

Value Added Services—Carrier provided enhancements to basic services.

Wirespeed—The speed of data transmission over a physical circuit.

Workgroup—A group of workstations and servers that commonly exchange data. This term is also used to describe a group of people who work together.

Workgroup Switching—The ability to handle asymmetric traffic patterns via high-speed (100 Mbps) interface and intelligent switching.

NOTES

NOTES

NOTES

Layer 3 Switching Components

FUNCTIONS:	INCLUDES:
Packet Switching	• TTL Decrement • MAC Address Rewrite • Recalculate Checksum
Route Processing	• Enhanced - IGRP • BGPv4 • STATIC • OSPF • PIM-SM • RIP • IGRP • PIM-DM • RIPv2
Intelligent Network Services	• Quality of Service (QoS) - IP Precedence - Weighted Round Robin Queueing - Weighted Random Early Detection • Hot Standby Router Protocol • Route Distribution • TACACS+ • Access Lists • Telnet, Ping, Traceroute • Network Debugging • DHCP Relay, Local Area Mobility • NTP • Syslog • and many more...

CISCO SYSTEMS

www.cisco.com
170 West Tasman, San Jose, CA 95134-1706
(408) 526-4000 • (800) 553-NETS (6387)
Fax (408) 526-4100

This Technology Guide is one in a comprehensive series of Guides that provide objective information and practical guidance on technologies related to Communications & Networking, the Internet, Document Management, Data Warehousing, and Enterprise Solutions. Our team of technical editors writes each Technology Guide to assist IT and business professionals in making informed decisions about all aspects of technology application development and strategic deployment.

techguide.com is supported by a consortium of leading technology providers. Cisco Systems has lent its support to produce this Guide.

Visit our Web Site at www.techguide.com to view and print this Guide, as well as all of our other Technology Guides. This is available as a free service.

techguide.com