

Tag Switching: Uniting Routing and Switching for Scalable, High-Performance Services

Abstract

The explosive growth of the Internet presents a serious challenge to service providers and equipment suppliers to keep up with this growth in traffic and the number of users. Service providers are also faced with the challenge of creating differentiated IP services and getting these new value-added services to market quickly. Cisco's Tag Switching technology is a key component in Cisco's plans for meeting these challenges. Tag Switching combines the performance and virtual-circuit capabilities of Layer 2 (data link) switching with the proven scalability and flexibility of Layer 3 (network) routing. Tag Switching is the first technology to fully integrate routing and switching for an Internet-scale environment, and it is one of several multilayer switching technologies being pioneered by Cisco. This white paper presents an overview of the benefits and technology of Tag Switching, and describes its application in service provider networks.

Evolving the Core of Service Provider Networks

Not a week goes by without a major publication running an article about the Internet and its rapid growth. Statistics show that the Internet is on a geometric growth curve, doubling the number of hosts and domains approximately once a year. Traffic growth is doubling at an even faster rate, fueled by graphic-rich Web pages and the beginnings of a multimedia explosion. For service providers, the challenge is to keep up with this growth in the core of their networks, while at the same time providing new services and capabilities to generate revenue and profits. For service providers that wish to offer services such as Frame Relay and Asynchronous Transfer Mode (ATM) along with value-added IP services and Internet service, the problem is even more complex. Compound this situation with the never-ending change that has characterized the data communications marketplace in recent years, and it's a wonder that managers of these networks sleep at all.

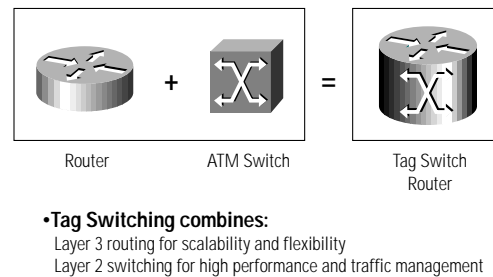
These trends lead to the following requirements for products and technologies for service provider networks:

- **Keep up with growth:** New products and technologies must support the rapid growth of network services, enabling service providers to build networks that scale in terms of traffic volume (both bits per second and packets per second), number of users, and number of nodes.
- **Reduce operations costs:** From laying new fiber to the manpower requirements of complex new technologies, service providers face rapidly escalating costs. New products and technologies for service provider networks should make efficient use of overtaxed bandwidth and reduce the complexity of operating the network.
- **Increase flexibility:** Service providers need an infrastructure that enables them to rapidly adjust their mix of services in response to changing market demands.
- **Increase reliability:** As electronic commerce increases and the size and complexity of the Internet continues to grow, providing highly reliable Internet and private IP services will become more challenging. New products and technologies must reduce this complexity and improve the robustness of service provider IP networks.
- **Generate new revenue from IP services:** The key to profitability in the IP era for service providers will be in providing differentiated, value-added services to their customers—not in providing commodity transport. New products and technologies should support rapid creation and deployment of such services.

Multilayer Switching

A key element in Cisco's overall strategy for providing scalable service provider network solutions is the integration of Layer 3 (network layer) routing, with its proven scalability of network size and its special services, and Layer 2 (data link layer) switching, with its high-performance and traffic-engineering capabilities. Cisco has developed a multilayer switching technology called Tag Switching that, in conjunction with other Cisco products, provides a leadership solution in meeting the requirements for network scalability and service delivery. This technology is specifically aimed at the requirements of large-scale wide-area networks (WANs).

Figure 1 Tag Switching: Uniting Routing and Switching for Scalable Services



Cisco is also developing multilayer switching technologies for enterprise LAN environments. In particular, NetFlow LAN Switching (NFLS) and Multiprotocol over ATM (MPOA) provide Layer 2/Layer 3 integration in LAN networks, with NFLS supporting Ethernet, Fast Ethernet, and Gigabit Ethernet backbones, while MPOA supports ATM backbones. Both are designed to provide high-performance, gigabit-capable solutions for multilayer switching in campus LAN environments.

Benefits of Tag Switching

Tag Switching provides benefits for service providers in three key areas:

1) IP+ATM Integration

Service providers understand that the key to profitability is value-added services. For data networking, IP is the protocol increasingly used by enterprise customers for their data, and is the place where value can be added by the service provider. However, support for IP quality of service (QoS) and IP multicast across an ATM backbone currently requires a complex and incomplete translation or mapping between the connectionless IP and connection-oriented ATM protocols. Tag Switching eliminates this translation loss and complexity by enabling ATM switches, such as the Cisco BPX[®] 8650, and 8680, Cisco MGX[™] 8800, and Cisco LightStream 1010, to directly support advanced IP services and protocols, reducing operational costs and bandwidth, and decreasing time to market for new services.

For example, IP class of service (CoS) will become available from many service providers during 1998. By providing two or three service classes (for example: premium, enhanced, standard), IP traffic that requires preferential treatment in the network when congestion occurs can be given priority service. IP CoS can be provided over an ATM core using ATM Forum QoS, but configuring

the many virtual circuits (VCs) and allocating bandwidth on a VC basis makes this a difficult network design. Tag Switching, by contrast, directly supports IP CoS, making it simple to configure the bandwidth allocation on a class-by-class basis for end-to-end IP QoS in a routed or ATM environment. Tag Switching leverages the QoS capabilities built into ATM switches such as the Cisco BPX series in providing IP QoS.

Tag Switching scales the Internet by integrating ATM with Layer 3 routing. Tag Switching enables ATM switches to be fully integrated into Internet networks, including Border Gateway Protocol (BGP) routing protocols, without the scalability problems of a pure Layer 2 ATM network ringed by a router overlay using permanent virtual circuits (PVCs) or switched virtual circuits (SVCs). In a Layer 2 environment, all routers on the edge of the ATM network need to peer with each other. In addition, the core network needs to cope with SVC reconnection storms if a link or switch fails. Both of these problems scale in proportion to the square of the number of routers creating very brittle networks with unmanageable failure characteristics.

With Tag Switching, each router sees fewer peers because all ATM tag switches participate in the Layer 3 protocols. Tag Switching further enhances scalability in an ATM network by reducing the number of tags needed in the network through a capability called "VC merge," which is similar to the concept of a multipoint-to-point VC.

Many service providers want to provide high-revenue services such as Frame Relay and cell relay access while simultaneously expanding their portfolio of high-growth IP services. Tag Switching provides the architectural flexibility to support ATM, Frame Relay, and IP services on the same high-performance ATM switched infrastructure. Furthermore, integrated WAN edge devices such as the Cisco MGX 8800 series can terminate all these services on a single platform. This scenario minimizes operational costs, accelerates time to market for new services, and increases customer loyalty.

2) Explicit Routing

Bringing Layer 2 benefits to Layer 3: Layer 2 technologies, such as ATM, make use of VCs, enabling these technologies to deliver certain capabilities that are currently difficult for connectionless routers to provide. Tag Switching brings the benefits of Layer 2 VCs to gigabit router backbones and integrated IP+ATM environments, including support for explicit routing and provisioning of IP VPN services.

By supporting explicit routing, Tag Switching has the capability to specify a path through a network at Layer 3 and to select a set of packets to follow that path. This scenario is accomplished by defining a filter that specifies the traffic and then mapping that filter to the appropriate TSP (Tag Switched Path). Several options exist for mapping a filter onto a TSP: mapping a filter to a specific TSP, mapping a filter to multiple TSPs, or load balancing between multiple TSPs. When mapping to multiple TSPs, the various TSPs can be ranked by attractiveness (for example, number of hops) with the default option of normal IP routing if no TSP is available.

The ability to control the flow of packets across a Layer 2 infrastructure to support load balancing has been one of the attractions of using ATM or Frame Relay switches in Internet cores. Explicit routing is important for four reasons:

- Fiber may not be available along the optimum path, based on where the traffic wants to flow.
- The cost of bandwidth may result in high-link utilization on some paths.
- A link failure may cause other links to become overloaded.
- Unexpected changes in traffic patterns can overload links.

Explicit routing enables the service provider to engineer the flow of packets to cope with these situations. Tag Switching extends these capabilities to IP networks that are built completely from Layer 3 routers, such as the Cisco 7500 router and the Cisco 12000 gigabit switch router (GSR). It also enables connectionless IP services built on an integrated IP and ATM network using switches, such as the BPX 8650 and MGX 8800, to be engineered at the IP layer without the complexities and service degradation that result from IP-to-VC translation.

3) Virtual Private Networks

VPNs for data are typically built from Frame Relay PVCs across a service provider's Frame Relay network. However, the growth of enterprises' intranets results in a traffic pattern shift from the traditional hub-and-spoke pattern (for which Frame Relay is a good match) to a mesh, any-to-any pattern that does not scale well with connection-oriented access services. Service providers are also encountering scalability limits in providing Frame Relay-based managed router service offerings. Tag Switching provides a means to support IP-based VPNs on either a Frame Relay/ATM backbone, integrated IP+ATM backbone, or a gigabit router backbone. By tagging packets based on both the destination address and on the VPN to which the packet belongs, Tag Switching-enabled products can support large numbers of IP VPNs with overlapping IP address spaces and end-to-end QoS-based service-level agreements (SLAs) with the same data security as a Layer 2 Frame Relay network.

Tag Switching enables IP VPNs to be built directly into the fabric of the network itself, making VPNs "visible" to the network. This capability is essential for:

- Adding value to the VPN (such as content-hosting services)
- Using the VPN as a foundation layer for delivering private value-added IP services
- Scalable deployments that use VPNs to form a large number of "closed user groups"

IP VPNs based on Tag Switching are connectionless, making an IP VPN a "virtual private internet." By providing the same networking paradigm as customer intranet LANs, integration with customer networks is seamless. Likewise, advanced IP services such as IP multicast are easily deployed in an IP VPN. Hybrid services that combine Frame Relay PVCs with IP VPN services can also be deployed over a single physical connection, generating current revenue while adding new high-growth IP services.

Based on the proven track records of Cisco in IP routing and ATM switching, Tag Switching is the only switching/routing integration technology that meets the key service provider requirements of scalability, advanced IP services, Layer 2 VC benefits, IP-based VPNs, the architectural flexibility to operate in conjunction with virtually any Layer 2 technology using any Layer 3 protocol, and standardization through the Internet Engineering Task Force's (IETF's) Multiprotocol Label Switching (MPLS)

working group. Other efforts to standardize the combination of Layer 2 and Layer 3, such as the ATM Forum's Multiprotocol over ATM, are not capable of scaling to the level required by service providers. Alternative Layer 2/Layer 3 combinations can only support one Layer 2 technology—ATM. Although Tag Switching formed the basis for MPLS, the final standard may differ in some respects from Cisco's current implementation of Tag Switching. Cisco is committed to full conformance with the MPLS standards as they are finalized.

Table 1 Benefits of Tag Switching

Capability	Benefit
IP+ATM Integration	<ul style="list-style-type: none">• IP services• Accelerates time to market for new services• No service degradation or translation loss from IP to ATM mapping• Reduces complexity for lower operational cost• Offers cost-effective integrated infrastructure• Provides scalable networks with high network robustness and availability• Maximizes return on investment by supporting ATM, Frame Relay, and Tag Switching on the same ATM infrastructure• Delivers hybrid IP and Frame Relay/cell relay services over a single connection
Explicit Routing	<ul style="list-style-type: none">• Brings Layer 2 VC functionality to a purely routed network• Leverages improved reliability through the creation of redundant routes• Offers higher return on investment through load balancing
Virtual Private Networks	<ul style="list-style-type: none">• Provides platform for rapid deployment of managed IP services, including intranets and extranets• Provides privacy for other value-added services such as voice, multimedia, and network commerce• VPN awareness in the network enables providers to participate as a peer with customer intranets, allowing for intranet application and content hosting services• Offers seamless integration with customer intranets• Offers increased scalability over current VPN implementations, with thousands of sites per VPN and hundreds of thousands of VPNs per service provider• Provides end-to-end IP QoS, with support for multiple classes of service• Offers easy management of VPN membership for rapid proliferation• Offers scalable any-to-any connectivity for extended intranets and extranets that encompass multiple businesses

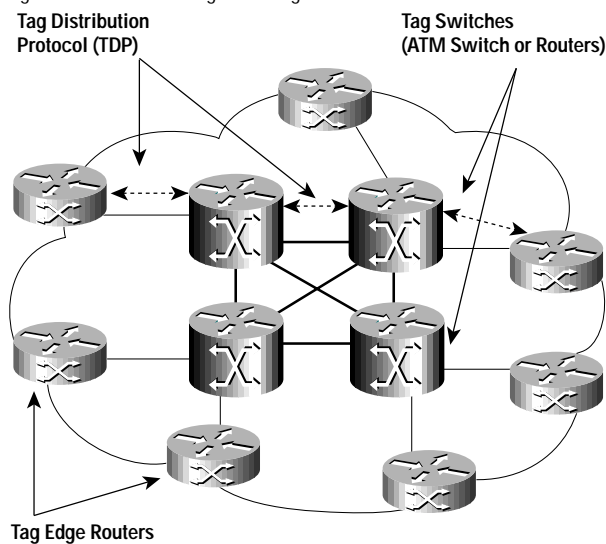
Tag Switching Overview

Tag Switching is an innovative new technique for high-performance packet forwarding that assigns "tags" to multiprotocol frames for transport across packet- or cell-based networks. It is based on the concept of "label swapping," in which units of data (for example, a packet or a cell) carry a short, fixed-length label that tells switching nodes how to process the data.

A Tag Switching internetwork consists of the following elements:

- Tag edge routers: Located at the boundaries of a service provider network, tag edge routers perform value-added network-layer services and apply tags to packets.
- Tag switches: Tag switches switch tagged packets or cells based on the tags. They may also support full Layer 3 routing or Layer 2 switching, in addition to Tag Switching.
- Tag Distribution Protocol (TDP): In conjunction with standard network-layer routing protocols, TDP is used to distribute tag information between devices in a Tag Switched network.

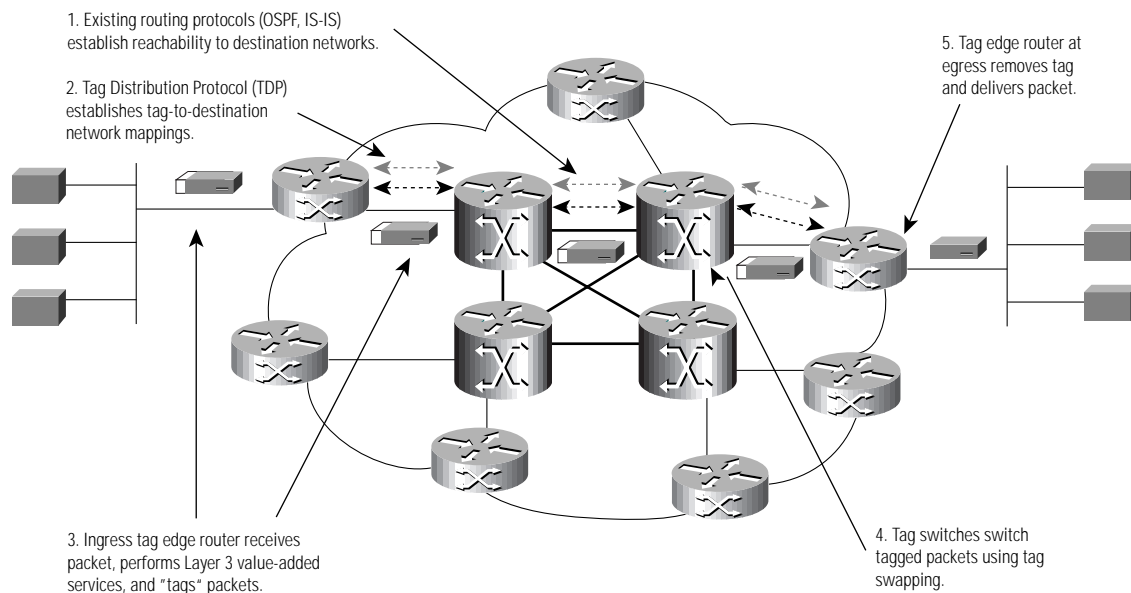
Figure 2 Elements of Tag Switching



The basic processing within a Tag Switching internetwork is as follows:

- 1) Tag edge routers and tag switches use standard routing protocols (for example, Open Shortest Path First [OSPF], Intermediate System-to-Intermediate System [IS-IS], and Enhanced Internet Group Management Protocol [EIGRP]) to identify routes through the network. These protocols fully interoperate with non-Tag Switching routers.
- 2) Tag switches use the tables generated by the standard routing protocols to assign and distribute tag information via TDP. Tag switches receive TDP information and build a forwarding database that makes use of the tags.
- 3) When a tag edge router receives a packet for forwarding across the tag network, it:
 - Analyzes the network-layer header
 - Performs applicable network-layer services
 - Selects a route for the packet from its routing tables
 - Applies a tag and forwards the packet to the next-hop tag switch
- 4) The tag switch receives the tagged packet and switches the packet based solely on the tag, without reanalyzing the network-layer header.
- 5) The packet reaches the tag edge router at the egress point of the network, where the tag is stripped off and the packet delivered.

Figure 3 Tag Switching Operation



Tag Edge Routers

Tag edge routers are full-function Layer 3 routing devices located at the edge of a tag switching network. A tag edge router can be a router, such as the Cisco 7500, or an IP+ATM WAN edge switch, such as the Cisco MGX 8800. Tags are applied to incoming packets, and tags are removed from outgoing packets. As full-function routers, tag edge routers also apply value-added Layer 3 services such as security, NetFlow accounting, and application-level QoS classification, and bandwidth management. Tag edge router capability does not require special hardware and is implemented as an additional set of software features within Cisco IOS[®] software, allowing existing installed routers to be software-upgraded with tag edge router capabilities.


A key part of the job of a tag edge router is to examine incoming packets and apply the proper tag to the packet before forwarding the packet. Tag Switching provides great flexibility in the algorithms and techniques it supports for mapping tags to packets. Examples include:

- **Destination prefix:** In this method, the tag edge router uses the normal router procedure of matching the destination IP address of the packet against the destination prefix entries in the forwarding tables of the router. This procedure determines the next hop for the packet. In a tag edge router, this lookup will also yield a tag value to apply to the packet. This technique allows traffic from multiple sources going to the same destination, to share the same tag (in addition, multiple destination prefixes can share the same tag value). Tag Switching greatly economizes on the number of tags required, relative to the number of source/destination IP address combinations seen by a router (this feature is one of the keys to the scalability of Tag Switching). This feature is particularly important in public Internets, where the number of source/destination address combinations can be enormous.
- **Class of service:** This method is similar to the destination prefix method, except that in this case, the IP Precedence field and the type-of-service (ToS) byte in the standard IP packet header is used to select the tag to use for the packet. This scenario enables different packets to be placed in differentiated IP service classes in a tag network.
- **Application flows:** This method looks at both the source and destination addresses, as well as other Layer 3 or Layer 4 information. This information can be used to provide finer granularity in processing the tagged packets and to maintain a given QoS through the network for a specific source/destination flow of packets.
- **Multicast:** This method uses the Protocol Independent Multicast (PIM) protocol to define the paths for multicast flows and to distribute the tags for multicast packets. Unique to the Tag Switching approach is the capability to enable ATM switches to directly support PIM and to become fully integrated into the IP infrastructure to support IP multicast services, even within an IP VPN.
- **Explicit routing:** Packets can be tagged so that they flow along specified routes, allowing network managers to balance the load across trunk circuits between nodes or cope with unbalanced network topologies during a node outage. This scenario is analogous to the use of VCs in pure Layer 2 devices for Internet traffic engineering, and can be implemented on Layer 3 routers using Tag Switching.
- **VPNs:** Tag Switching can also be used to support multiple overlapping IP address spaces on a common IP infrastructure. By tagging a packet based on both its destination IP address and its VPN, multiple private IP networks can be accommodated, and customers can subscribe to multiple VPNs.

Tag Switches

Tag switches are the core of a Tag Switching internetwork. Tags are short, fixed-length labels, enabling IP+ATM switches to do simple and fast table lookups. This scenario enables tag switches to implement the lookup and forwarding capabilities using fast hardware techniques, including ATM cell switching.

Because Tag Switching decouples the tag distribution mechanisms from the data flows, it supports a wide range of physical and link-layer technologies. For ATM, the tag is placed in the ATM cell header in the virtual path identifier/virtual channel identifier (VPI/VCI) field. For use with the Point-to-Point Protocol (PPP), the tag is placed between the Layer 2 and Layer 3 headers. In a LAN, the tag is placed after the MAC header. This setup enables Tag Switching to be used over a wide variety of media, including ATM links, packet-over-SONET links, Ethernet, and so forth. Tag Switching also is not specific to IP, because the routing protocols are separate and standard Tag Switching can be used to support multiple Layer 3 protocols (for example, IPv4, IPv6).



When outfitted with Tag Switching software, standard routers can act as tag switches. By supporting TDP, and adding the capability to switch tagged packets based on the tag values, Internet core routers can participate in a tag network backbone. For current routers, this scenario brings the explicit routing and IP VPN capabilities of Tag Switching to pure router networks, a significant enhancement over today's capabilities. Separate Layer 2 switching backbones are not required to gain these traffic engineering benefits.

IP+ATM switches, such as the Cisco BPX 8650, and the Cisco LightStream 1010, can also be used as tag switches. These switches implement the appropriate, standard Layer 3 routing protocols as well as TDP. Tags are placed in the VPI/VCI fields of cells by the tag edge routers, and the ATM switch switches cells based on the VPI/VCI values, as it does for all cells.

A key difference between Tag Switching and ATM is that standard ATM uses a connection setup procedure to allocate virtual circuits and program the ATM switching hardware, whereas Tag Switching uses standard routing protocols and TDP. The result is that ATM switches that perform Tag Switching have no call setup overhead for IP traffic.

Because Tag Switching and ATM Forum-compliant ATM can coexist on the same ATM switch, Cisco's IP+ATM switches can provide both Internet service and ATM services on the same platform. Tag Switching avoids the use of SVCs for highly dynamic IP packet flows and frees CPU processing power for Private Network-Network Interface (PNNI) and the needs of longer-lived ATM VCs, such as real-time voice or video flows.

Tag Distribution Protocol

TDP provides the means by which tag switches exchange tag information with other tag switches and with tag edge routers. The tag edge routers and tag switches build their routing databases using standard routing protocols (for example, OSPF, IS-IS, EIGRP). Neighboring tag switches and edge routers then distribute tag values to each other using TDP, for storage in a tag forwarding information base (TFIB). Unlike standard ATM, there is no call setup procedure.


When the destination-prefix tagging algorithm is combined with TDP and standard routing protocols, tagging information in the TFIB is established before packet flows traverse the network. This topology-driven approach means that all packets, even packets on short-lived flows, can be tag switched. Some IP switching implementations use a traffic-driven approach that sets up individual flows only after some number of packets with the same source/destination have passed by. In a high-performance switch based on ATM, sending many packets off to a Layer 3 function separate from the ATM cell path causes a degradation of performance and QoS and can be a significant bottleneck in scaling such IP switching implementations. In contrast, the Tag Switching destination prefix method provides higher performance by switching all packets at the tag level.

Another advantage of using standard routing protocols and TDP in a tag switched network is seen when an IP network is built from a core of ATM switches surrounded by edge routers. In this configuration today, the edge routers are mesh connected by VCs through the ATM switches, meaning that the edge routers are all peers of each other from a routing protocol perspective. This topology scales only up to a limited number of routers before the large number of peers causes problems with the routing protocols. The problem is that the ATM switches are transparent to the routers. In a tag switched network, the ATM switches, acting as tag switches, participate fully in the hierarchical routing protocols and act as peers to the tag routers at the edges. Therefore, the tag edge routers see far fewer peers, and hence the size of the network, measured in number of routers, can scale to much larger sizes.

Application in Multiservice Networks

In networks that provide both IP service and ATM/Frame Relay services on a common core ATM infrastructure, Tag Switching provides an ideal way to integrate the ATM platforms into a high-performance IP service offering. Cisco's IP+ATM switches simultaneously support IP services via Tag Switching and standard ATM switching using ATM Forum and other industry-standard protocols. Tag Switching enables a closer coupling of the ATM switch into the IP service, providing a stronger reason for using a common ATM core for all service offerings. Using an integrated IP+ATM WAN edge switch such as the Cisco MGX 8800, IP, Frame Relay and ATM services, including IP VPNs, can be provisioned over a single physical link with separate PVCs. This setup offers improved service to the customer, significantly lowers capital and operational costs, and accelerates time to market for new services. It also gives the service provider great flexibility in adjusting the mix of provisioned services in response to changing market demands.

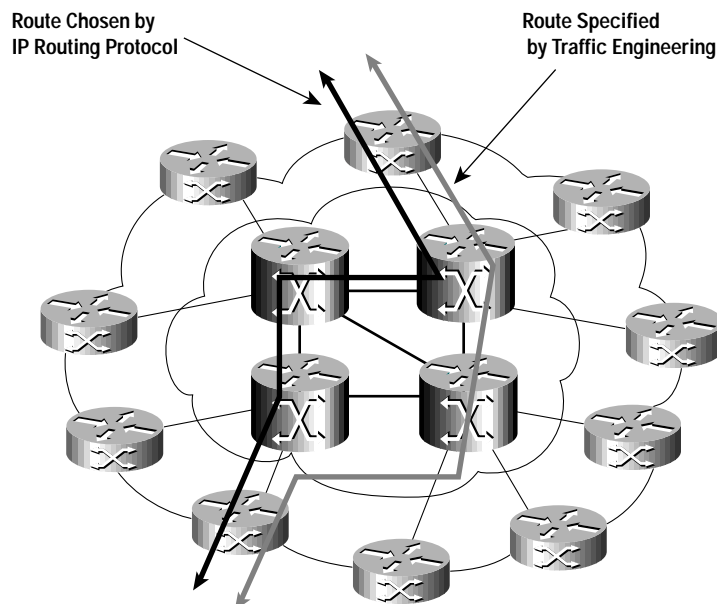
Figure 4 Tag Switching in a Multiservice Network



Application in Router-Only Internets

Because Tag Switching can be used on standard routers, a network built exclusively from routers can deploy Tag Switching software to gain the advantages of traffic engineering and IP-based VPNs provided by Tag Switching. For high performance in this environment, Tag Switching can support packet over SONET OC-3 interfaces on core routers. Tag Switching will also be supported on the Cisco 12000 GSR to enable multigigabit speeds to OC-12, OC-48, and beyond.

Figure 5 Tag Switching in a Gigabit Router Internet



Network Evolution

Together, the rate of growth and the need for extended use of assets to become profitable makes it unrealistic for service providers to start from scratch with a new technology. For these reasons, Cisco's implementation of Tag Switching has been developed to allow rapid evolution of a network toward a tag switched infrastructure, but at the same time to protect existing and future investments.

For a pure router network, the evolution will be straightforward. The existing Cisco 7500- or 12000-based infrastructure will first receive a software upgrade to a version of Cisco IOS software that includes both Tag Switching and routing. TDP will be enabled, but at the same time, the tag switch routers can be configured to prefer the normal routing path. When TDP has been successfully deployed, the network can be progressively cut over to the Tag Switching approach. No new hardware or major disruption of service will be required. Note that traffic engineering can be enabled along with normal IP routing for most packets (that is, only the packets that take a different path are tagged; all other packets are routed as normal). This setup provides another option for introducing Tag Switching into a network in a controlled manner.

For a multiservice network based on the Cisco BPX and the Cisco 7500, a similar approach will be taken. A network of Cisco BPX 8650s can be built to support many services, including IP. The IP services would today be supported by a Cisco 7500-based router overlay that provides IP over ATM facilities. To cut over to a Tag Switching approach, the network will require the following upgrade process:

- A Cisco BPX software upgrade to a release that will allow external Tag Switching control; this upgrade can be done without significant traffic interruption
- A Cisco 7500 or 7200 software upgrade to a Cisco IOS release that includes Tag Switching and Cisco BPX control
- Establishment of the tag switching network through switch and router configuration
- Progressive cutover of traffic to the tag switched network

Again, no new hardware or major disruption of service is required. For a network based on the Cisco LightStream 1010, a similar migration approach can be used.

Cisco's new MGX 8800 wide-area service provider edge switch has integrated IP+ATM capabilities, with built-in routing and Cisco IOS software. The MGX 8800 fully interoperates with other tag edge routers, making it an ideal choice for new deployments. Service providers can offer Frame Relay, private IP, Internet, and circuit emulation services using the same edge device, with the flexibility to determine the right mix of services as business needs dictate.

The approach being taken in deploying Tag Switching underscores Cisco's commitment to investment protection and to the continual nondisruptive growth of IP and ATM/Frame Relay networks.

Tag Switching and Related Technologies

Several related technologies have been proposed by other vendors for Layer 2/Layer 3 integration in the Internet, including Newbridge's Carrier Scale Internetworking (CSI), Ascend's IP Navigator, and Nokia/Ipsilon's IP Switching.

Newbridge's CSI architecture leverages MPOA to provide a way to support IP across a standard ATM backbone. CSI uses edge devices based on Newbridge's VIVID LAN products and on separate route servers to provide the routing functionality in the network. Because MPOA, as a standard, is designed for LAN environments and because MPOA has some VC scalability issues, Newbridge is modifying MPOA to create the proprietary CSI protocol set (for example, removed support for LAN Emulation, added multipoint-to-point VC support, and so forth). The key concerns expressed within the industry regarding this approach include:

- Will the routing information distribution method used by MPOA and CSI scale to the size of networks that form the Internet (the standard version of this protocol is called the Next-Hop Resolution Protocol [NHRP]).
- CSI is a closed, proprietary architecture (in other words, it differs from MPOA), and Newbridge has not publicly disclosed the technical details of CSI or its protocols.

- The CSI architecture does not have a sensible way of allocating bandwidth to particular IP services and will require inefficient use of network resources. The CSI approach is ATM-centric and does not implement an end-to-end IP model, as is provided by Tag Switching.

Ascend's IP Navigator technology is also a Layer 2/Layer 3 integration technique. Like CSI, it uses a proprietary protocol (Virtual Network Navigator [VNN]) within the IP Navigator network (hence all components within the network must be a limited set of Ascend products). It is targeted at providing IP-based VPN capabilities on an Ascend ATM backbone, not at supporting public Internet service. Like CSI, it supports ATM only in the backbone (the Ascend GRF router does not support IP Navigator). It should be noted that although Cascade is active in the MPLS working group, to date it has stated support only for MPLS at the edge of an IP Navigator network.

Nokia/Ipsilon's IP switching is another technology for combining ATM Layer 2 switching with Layer 3 routing. IP switching typically allocates a label per source/destination packet flow. An IP switch processes the initial packets of a flow by passing them to a standard router module that is part of the IP switch. When an IP switch has seen enough packets go by on a flow to consider it long-lived, the IP switch sets up labels for the flow with its neighboring IP switches or edge routers such that subsequent packets for the flow can be label-switched at high speed (for example, in an ATM switching fabric), bypassing the slower router module. As noted previously, there is a general industry consensus that host-to-host flow approaches will not scale to the level required in large IP networks or the core of the Internet.

Of currently proposed multilayer switching technologies, Tag Switching provides the best match with the requirements for such a technology. The Tag Switching topology-driven approach, which couples its destination prefix algorithm to standard routing protocols, supports much more efficient use of labels than traffic-driven per-flow schemes, and avoids flow-by-flow setup procedures altogether. Tag Switching provides direct support for advanced IP services, such as CoS, RSVP, IP VPNs, and multicast on ATM switches, and brings the benefits of explicit routing and VPNs to gigabit routers. Most importantly, Cisco's implementation of Tag Switching will conform to the MPLS standard, both for external interfaces and for interfaces within the core of the network, enabling multivendor networks.

The inherent flexibility of Tag Switching also provides an outstanding match with the evolutionary requirements of public and private IP networks. Tag Switching is designed from the ground up to support both packet and cell interfaces. As a result, Tag Switching can be implemented as a software upgrade on existing Internet Cisco 7500 routers, providing traffic engineering benefits and a smooth migration from today's Internet to future tag switched infrastructures. Tag Switching on the Cisco BPX 8650 and 8680, MGX 8800, and the LightStream 1010 realize the full benefits of integrating ATM speeds with Cisco's proven and powerful Cisco IOS routing software, while providing full support for standards-compliant ATM software on the same platform. Tag edge router functionality will also become part of Cisco IOS software, and will be available as a software upgrade on existing installed Cisco routers and BPX ATM switches.

Cisco's Plans for Tag Switching

Cisco has submitted the TDP to the IETF MPLS working group for standardization and has published extensive Internet drafts that describe Tag Switching. Cisco has licensed Cisco IOS Tag Switching technology to partners, such as Alcatel, and has announced successful initial MPLS interoperability testing.

Summary

Tag Switching is a new technology developed from the routing and switching expertise of Cisco Systems. It provides scalability through integrated support of ATM switches within an IP core, supports advanced IP services on ATM switches, and delivers traffic engineering and IP-based VPN capabilities on both ATM switches and on standard Layer 3 routers. Through support for the MPLS standard as it evolves, Tag Switching enables true multivendor networking. As one of Cisco's multilayer switching technologies, Tag Switching is a breakthrough technology for service providers that will support scaling of IP services and the Internet well into the next century.



Corporate Headquarters

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

European Headquarters

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

Americas

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

Asia Headquarters

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

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

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland • Singapore