Scaling the Internet With Tag Switching

White Paper

- **News Release: Cisco Scales the Internet**
- **News Release: Cisco's New Tag Switching Technology**
- **News Release: Industry's First ATM/IP Switching Wide-Area System**

- **Internet Draft: Tag Switching Architecture Overview**
- **Internet Draft: Tag Distribution Protocol**

Abstract

The explosive growth of the Internet presents a serious challenge to Internet service providers and equipment suppliers in keeping up with this growth in traffic and number of users. Cisco's new Tag Switching technology is a key component in Cisco's plans for meeting this challenge. Leveraging the strengths of Cisco and StrataCom, tag switching combines the performance and traffic management capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. Tag switching is the first technology to fully integrate routing and switching for an Internet-scale environment, and 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.

Scaling the Core of the Internet

Not a week goes by without a major publication running an article on the Internet and its rapid growth. Statistics bear this out, showing 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 graphics-rich web pages and the beginnings of a multimedia explosion. For Internet 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 ATM along with Internet service, the problem is even more complex. Compound this 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 key requirements for products and technologies for the core of the Internet:

- **Capacity:** the ability to provide the forwarding and routing capacity and performance for public Internet services
- **Scalability:** the ability to grow the size of the network along with the growth in number of users and amount of traffic
- **Traffic Management:** the ability to efficiently utilize wide area network links between core Internet sites
- **Service Differentiation:** the ability to support new Internet services, such as Quality of Service (QoS)
- **Evolutionary:** the ability to adapt to the changing requirements of Internet networks, while avoiding fork lift upgrades

Integration of Routing and Switching

A key element in Cisco's overall strategy for providing scaleable Internet 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. By combining the industry-leading expertise of Cisco and StrataCom in these areas, Cisco has developed a technology called "tag switching" that, in conjunction with other Cisco products, provides a leadership solution in meeting the requirements for Internet core scalability.

Tag switching provides three key benefits for service providers:

- **Scalability:** Tag switching scales the Internet by integrating ATM with Layer 3 routing. Tag switching enables ATM switches, such as the Cisco StrataCom BPX, to be fully integrated into Internet core networks, including BGP routing protocols, without the scalability problems of a pure Layer 2 network ringed by a router overlay. Internets can take advantage of the high performance of the Cisco StrataCom BPX, as well as the ability to provide Internet and ATM/frame relay services on the same platform.
- **Traffic Management:** Tag switching simplifies traffic management in router-based Internets by integrating Layer 2 circuit capabilities. 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. Tag switching supports this capability, but more importantly, extends these capabilities to Internets that are built completely from Layer 3 routers. This increases the control that network managers have over the flow of packets in router-based Internets.
- **Performance:** Tag switching enables higher performance platforms by simplifying packet forwarding and switching decisions. This simplification will allow future ATM switches and router platforms to support multi-gigabit, high speed interfaces. Future platforms can employ special tag switching hardware to gain performance advantages from the simplified nature of packet forwarding under tag switching.

Based on the proven track records of Cisco and StrataCom in Layer 3 routing and Layer 2 switching, tag switching is the only switching/routing integration technology that meets the key core Internet requirements of capacity, scalability, traffic management, and services support in an evolutionary and flexible manner.

Tag Switching Overview

Tag switching is a 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 (e.g., 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 an Internet, tag edge routers perform value-added network layer services and apply tags to packets.
- **Tag switches:** switch tagged packets or cells based on the tags. Tag switches 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 Internet.

Figure 1.: Tag Switching Internet

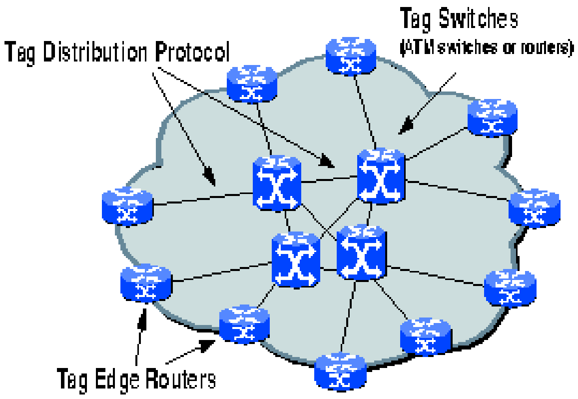

The basic processing within a tag switching internetwork is as follows:

1. Tag edge routers and tag switches use standard routing protocols (e.g., EIGRP, BGP, OSPF) to identify routes through the network. These fully interoperate with non-tag switching routers.
2. Tag routers and switches use the tables generated by the standard routing protocols to assign and distribute tag information via the tag distribution protocol (TDP). Tag routers receive the TDP information and build a forwarding database which 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 re-analyzing 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.

Tag Edge Routers

Tag edge routers are full-function Layer 3 routing devices located at the edge of a tag switching network. They apply tags to incoming packets and remove tags from outgoing packets. As full-function routers, tag edge routers also apply value-added Layer 3 services, such as security, accounting, and QoS classification. Tag edge router capability does not require special hardware and is implemented as an additional set of software features under Cisco IOS, allowing existing installed routers to be software-upgraded with tag edge router capabilities.

A key part of a tag edge router's job 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 packet's destination IP address against the destination prefix entries in the router's forwarding tables, which 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), and avoids the "label explosion" problem of current IP switching implementations. Tag switching greatly economizes on the number of tags required, relative to the number of source/destination IP address combinations seen by a router, and is a key to the scalability of tag switching. This is particularly important in public Internets, where the number of source/destination address combinations can be enormous.
- **Traffic tuning:** Packets can be tagged such 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 is analogous to the use of virtual circuits in pure Layer 2 devices for Internet traffic engineering, but, as will be discussed below, can be implemented on Layer 3 routers using tag switching.
- **Application flows:** This method looks at both the source and destination address, as well as other Layer 3 information. This can be used to provide finer granularity in processing the tagged packets and maintain a given quality of service through the network for a specific source/destination flow of packets, such as for RSVP.

With its flexibility of mapping techniques, tag switching is complementary to Cisco's NetFlow switching. NetFlow switching provides high performance processing of value-added Layer 3 services in tag edge routers at the edge of a network. Tag switching provides high performance switching of packets in the core of the network. NetFlow-switched packets can be tagged in the tag edge router via any of the means above, depending on the specific processing required for the packet through the network (i.e., packets that only need special Layer 3 processing at the edge can be tagged via destination prefix, while packets that require special QoS processing within the network can use flow-based tagging).

Tag Switches

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

Since tag switching decouples the tag distribution mechanisms from the data flows, a wide variety of methods of associating a tag with a packet can be used and will interoperate in a tag network, including:

- In the Layer 2 header (e.g., in the VCI field for ATM cells)
- In the Layer 3 header (e.g., in the flow label field in IPv6)
- In between the Layer 2 and Layer 3 headers

This enables tag switching to be used over a wide variety of media, including ATM links, Packet-over-SONET (POS) links, ethernet, etc. Tag switching is also not specific to IP. Since the routing protocols are separate, and are standard, tag switching can be used to support multiple Layer 3 protocols.

When outfitted with tag switching software, standard routers can act as tag switches. By supporting the tag distribution protocol, and adding the ability to switch tagged packets based on the tag values, Internet core routers can participate in a tag network backbone. For current routers, this brings the traffic tuning capabilities of tag switching to pure router Internets, a significant enhancement over today's capabilities. No separate Layer 2 switching backbone is required to gain these traffic tuning benefits. In the future, special tag processing hardware in tag switching routers can boost the performance of tag packet processing to unprecedented levels.

ATM switches, such as the Cisco StrataCom BPX, can also be used as tag switches. To act as a tag switch, an ATM switch will implement the appropriate, standard Layer 3 routing protocols, as well as the tag distribution protocol. Tags will be placed in the VCI fields of cells by the tag edge routers, and the ATM switch will switch cells based on the 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 VCI's and program the ATM switching hardware, while tag switching uses standard routing protocols and the tag distribution protocol. The result is that ATM switches performing tag switching do not need to handle high call setup rates.

Since tag switching and ATM Forum-complaint ATM can coexist on the same ATM switch, platforms like the Cisco StrataCom BPX can provide both Internet service and ATM services on the same platform. Tag switching avoids the use of switched virtual circuits (SVC's) for highly dynamic IP packet flows and frees CPU processing power for P-NNI and the needs of longer-lived ATM virtual circuits, such as real-time voice or video flows.

Tag Distribution Protocol

The 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 (e.g., EIGRP, BGP, OSPF). Neighboring tag switches and edge routers then distribute tag values to each other using TDP, for storage in a tag information database (TIB). 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 TIB is established prior to packet flows traversing the network. This means that all packets of a flow of packets can be tag switched, and that even packets on short-lived flows can be tag switched. Some IP switching implementations only set up a switched flow after some number of packets with the same source/destination pass 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 can be a significant bottleneck in scaling such IP switching implementations. In contrast, tag switching's destination prefix method provides higher performance by switching more packets at the tag level.

Another advantage of using standard routing protocols and TDP in a tag switched network is seen when an IP Internet is built from a core of ATM switches surrounded by edge routers. In this configuration today, the edge routers are mesh connected by VC's through the ATM switches, which means that the edge routers are all peers of each other, from a routing protocol perspective. This topology will scale 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, and not routing peers. 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. This means that 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 Router-Only Internets

Since tag switching can be used on standard routers, an Internet built exclusively from routers can deploy tag switching software to gain the advantages of traffic tuning provided by tag switching. For high performance in this environment, tag switching can support Packet Over SONET (POS) OC-3 interfaces on core routers. Tag switching is also extensible to future, higher performance router platforms and can enable multi-gigabit speeds to OC-12, OC-48 and beyond.

Figure 2.: Tag Switching in a Router-Only Internet

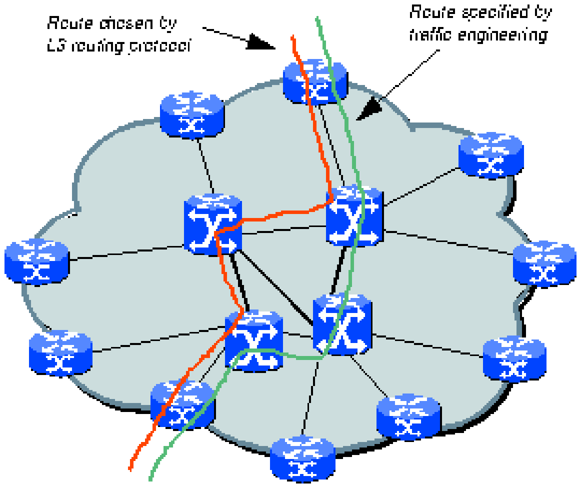

Application in Multiservice Networks

In networks that provide both Internet 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 Internet service offering. It enables an ATM platform, like the BPX, to simultaneously support Internet core switching 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 Internet service, providing a stronger reason for using a common ATM core for all service offerings.

Figure 3. : Tag Switching in a Multiservice Network

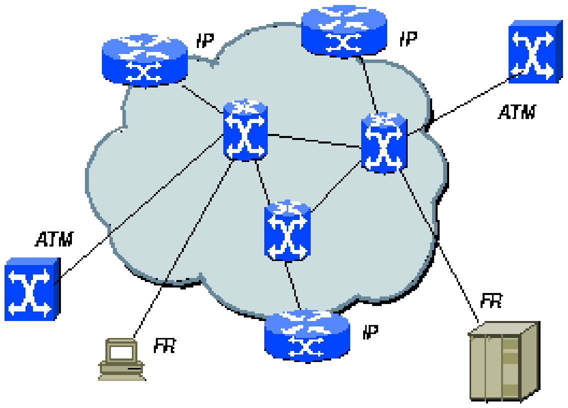

Network Evolution

Together, the rate of growth of the Internet and the need for extended use of assets to become profitable makes it not viable for service providers to start from scratch with a new technology. For these reasons, Cisco’s implementation of tag switching has been developed in a way 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-based infrastructure will first receive a software upgrade to a version of Cisco IOS which includes both tag switching and routing. The TDP protocol will be enabled, but at the same time, the tag switching routers can be configured to prefer the normal routing path. Once TDP has been successfully deployed, the network can be progressively cutover to the tag switching approach. No new hardware or major disruption of service will be required.

For a multi service network based on Cisco StrataCom BPX’s and Cisco 7500’s, a similar approach will be taken. A network of Cisco StrataCom BPX’s can be built to support many services including IP. The IP services would today be supported by a Cisco 7500-based router overlay which would provide IP over ATM facilities. To cutover to a tag switching approach, the network will require the following upgrade process:

1. A Cisco StrataCom BPX software upgrade to a release which will allow external tag switching control; this can be done without significant traffic interruption
2. A Cisco 7500 software upgrade to a release which will include tag switching and Cisco StrataCom BPX control
3. Establishment of the tag switching network through switch and router configuration
4. Progressive cutover of traffic to the tag switched path

Again, no new hardware or major disruption of service will be required.

The approach being taken to deploying tag switching underscores Cisco's commitment to investment protection and to the continual non-disruptive growth of the Internet.

Tag Switching and IP Switching

IP switching is a related technology for combining ATM Layer 2 switching with Layer 3 routing (i.e., IP switching is also a type of multilayer switching). 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 (e.g., in an ATM switching fabric), bypassing the slower router module. Special IP switching gateways are responsible for converting packets from non-labeled to labeled format, and from packet media to ATM.

Tag switching greatly enhances upon the label swapping concept, resulting in a highly scalable and flexible technology. Tag switching's destination prefix algorithm, coupled with standard routing protocols, supports much more efficient use of labels than per-flow schemes, and avoids flow-by-flow setup procedures altogether. This yields the scalability required for public Internet service networks, where the number of flows is enormous and the rate of change of flows is very high. By pre-establishing tag mappings at the same time as routing tables are populated, tag switching can tag switch both short lived flows and the initial packets of long-lived flows, avoiding bottlenecks in high performance applications.

Tag switching's inherent flexibility also provides an outstanding match with the evolutionary requirements of public internets. 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 tuning benefits and a smooth migration from today's Internet to future tag switched infrastructures. Tag switching on the Cisco Stratacom BPX will 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 and be available as a software upgrade on existing installed Cisco routers, avoiding the need to do a hardware upgrade to special-purpose, limited function IP switching gateways.

Cisco's Plans for Tag Switching

Cisco is actively working with a large group of vendors to establish an IETF working group to standardize tag switching. Cisco plans to license any patents it receives related to the tag switching standards on open terms for no (i.e., $0) license fee. The company also plans to license Cisco's implementation of tag switching in Cisco IOS to its Cisco IOS partners.

Tag switching technology will be available for field trials on the Cisco 7500 series routers in the first half of 1997. Versions for the Cisco LightStream 1010 ATM Switch and StrataCom BPX will be available later in 1997.

Tag switching technology is also relevant in large campus switched router environments. Running routing protocols within existing switching systems --- for instance, ATM switches --- allows these switches to participate in routing protocols with conventional routers. This can greatly increase network scalability by reducing the number of routing peers that edge routers must deal with, and avoiding the well-known problems of large flat Layer 2 networks in the backbone. Tag switching is complimentary to ATM LAN technologies such as LANE and MPOA, which focus on providing VLAN and inter-VLAN cut-through solutions respectively.

Summary

Tag switching is a new technology developed from the routing expertise of Cisco and the switching expertise of StrataCom. It provides scalability through integrated support of ATM switches within an Internet core, traffic tuning capabilities on both ATM switches and on standard Layer 3 routers, and enables future higher performance implementations via its simple forwarding structure. As one of Cisco’s multilayer switching technologies, tag switching is a breakthrough technology for service providers that will support scaling of the Internet well into the next century.

Posted: Thu Oct 9 15:16:06 PDT 1997