

# Tag Switching Overview

---

Tag Switching combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. It enables service providers to meet challenges brought about by explosive growth and provides the opportunity for differentiated services without necessitating the sacrifice of existing infrastructure. The Tag Switching architecture is remarkable for its flexibility. Data can be transferred over any combination of Layer 2 technologies, support is offered for all Layer 3 protocols, and scaling is possible well beyond anything offered in today's networks.

Specifically Tag Switching can efficiently enable the delivery of IP services over an ATM switched network. It supports the creation of different routes between a source and a destination on a purely router-based Internet backbone. Service providers who use Tag Switching can save money and increase revenue and productivity.

Tag Switching offers the following benefits:

- IP over ATM scalability—Enables service providers to keep up with Internet growth
- IP services over ATM—Brings Layer 2 benefits to Layer 3, such as traffic engineering capability
- Standards—Supports multivendor solutions
- Architectural flexibility—Offers choice of ATM or router technology, or a mix of both

## Tag Functions

In conventional Layer 3 forwarding, as a packet traverses the network, each router extracts all the information relevant to forwarding the packet from the Layer 3 header. This information is then used as an index for a routing table lookup to determine the packet's next hop.

In the most common case, the only relevant field in the header is the destination address field, but in some cases other header fields may also be relevant. As a result, the header analysis must be done independently at each router through which the packet passes, and a complicated lookup must also be done at each router.

In Tag Switching, the analysis of the Layer 3 header is done just once. The Layer 3 header is then mapped into a fixed length, unstructured value called a *tag*.

Many different headers can map to the same tag, as long as those headers always result in the same choice of next hop. In effect, a tag represents a *forwarding equivalence class*—that is, a set of packets, which, however different they may be, are indistinguishable to the forwarding function.

The initial choice of tag need not be based exclusively on the contents of the Layer 3 header; it can also be based on policy. This allows forwarding decisions at subsequent hops to be based on policy as well.

Once a tag is chosen, a short tag header is put at the front of the Layer 3 packet, so that the tag value can be carried across the network with the packet. At each subsequent hop, the forwarding decision can be made simply by looking up the tag. There is no need to re-analyze the header. Since the tag is a fixed length and unstructured value, looking it up is fast and simple.

## Distribution of Tag Bindings

Each *tag switching router* (TSR) makes an independent, local decision as to which tag value is used to represent which forwarding equivalence class. This association is known as a tag binding. Each TSR informs its neighbors of the tag bindings it has made. This is done by means of the Tag Distribution Protocol (TDP).

When a tagged packet is being sent from TSR A to a neighboring TSR B, the tag value carried by the packet is the tag value that B assigned to represent the packet's forwarding equivalence class. Thus the tag value changes as the packet travels through the network.

## Tag Switching and Routing

A tag represents a forwarding equivalence class, but it does not represent a particular path through the network. In general, the path through the network continues to be chosen by the existing Layer 3 routing algorithms such as OSPF, Enhanced IGRP, and BGP. That is, at each hop when a tag is looked up, the next hop chosen is determined by the dynamic routing algorithm.

## Tag Switching and Traffic Engineering

In conventional Layer 3 routing, network topologies frequently include multiple paths between two points, but the normal routing procedure is to select a single path as the Layer 3 route between two points regardless of the load on the links that implement the path. As a consequence, some links are congested and some are underused.

*Traffic engineering* provides a way to override routing protocols across multiple routers. It gives you the ability to direct selected traffic over specific paths in the network in order to efficiently use network resources and provide different levels of service.

To engineer your network traffic, you follow a two-step process. First, you define a sequence of links between two routers. Tag switching is used to tunnel packets between the two routers over these links. The links collectively form a *tag switched path* (TSP) *tunnel*, which defines a traffic engineering path. Second, you select the traffic which you want forwarded on to the tunnel.

## Traffic Engineering Tunnels and Filters

The traffic to be engineered is specified by a traffic engineering filter. The filter is associated with a TSP tunnel using a traffic engineering path.

The router at the head of the tunnel arranges that packets that match the filter be injected into the tunnel rather than being forwarded to their Layer 3 next hop. Injection consists simply of sending the packet to the first hop in the tunnel with a tag that causes that first hop to send the packet to the second hop of the tunnel, and so on.

For the initial release of traffic engineering, the only supported filtering is by "egress address." This filter matches traffic whose destination or BGP next hop is "address."

Multiple tunnels with different preferences can be specified for a single filter. A preference is an option you can select among multiple candidate routes for a filter, with the lower-valued preference being more desirable. The most preferred of the acceptable tunnels is used for the traffic.

A loop prevention algorithm operates to ensure that a tunnel is not used for traffic that might loop back to the head of the tunnel.

## Traffic Engineering Tunnel Configuration

Configuration and the initiation of the tunnel are controlled by the *headend* (transmit end) router. Per-tunnel configuration of other routers is unnecessary.

Routers create and maintain the *traffic engineering tunnels* based on information you enter through the command line interface (CLI). See the chapter “Tag Switching Commands” in the *Cisco IOS Switching Services Command Reference*.

