

Feature Brief

# Integration with Multi-protocol Label Switching (MPLS)

---



Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
408 745 2000 or 888 JUNIPER  
[www.juniper.net](http://www.juniper.net)

Part Number: 210005-001 Aug. 2005

## Overview

The Juniper Networks WX™ and WXC™ application acceleration platforms, two members of a larger family of solutions that improve application response times, provide IT with a range of technologies to overcome the technical limitations of WANs and speed application delivery across wide-area links. These features are implemented as part of the integrated WX Framework, which defines specific attributes that an application acceleration platform must have to overcome the bandwidth, latency, congestion and manageability issues that impede application performance over the WAN. Each element of the WX Framework addresses a specific challenge that prevents applications from running efficiently over the WAN. Those elements are organized into the following four categories:

### Compression and Caching:

The WX Framework includes Molecular Sequence Reduction™ (MSR™) technology, next-generation, memory-based compression that frees up WAN capacity by eliminating repeated data patterns. The MSR feature is complemented by the Network Sequence Caching technology, which uses hard disks to store and recognize large repeated patterns, even if they were sent days or weeks earlier.

### Acceleration:

The acceleration component of the WX Framework includes Packet Flow Acceleration™ (PFA™) techniques, which combat the effects of latency on the TCP protocol. The Application Flow Acceleration™ (AppFlow™) technology augments that TCP acceleration with protocol-specific acceleration for applications such as Exchange, Microsoft file services, and web. The AppFlow feature pipelines multiple data blocks and web objects across the WAN, improving user productivity by reducing their wait times as multiple round trips complete.


### Application Control:

The WX Framework includes Policy-based Multipath™ (Multipath™) optimization technology, which allows IT to direct specific application flows to a specific WAN link when multiple links are available. Application control features also include Quality of Service (QoS) and bandwidth-management tools for prioritizing critical application traffic and ensuring bandwidth availability.

### Visibility:

The visibility components of the WX Framework include WebView device management for configuring and managing individual WX and WXC application acceleration platforms, and the WX Central Management System™ (CMS™) software for gaining visibility into and centralized control over WX and WXC platforms distributed throughout an organization.

**The Integrated WX Framework**



Each of these elements interacts with one another to dynamically adjust and improve their capabilities. Working in concert, these features provide IT with the greatest degree of application acceleration and WAN optimization.

One of the most critical issues to address in WAN transmissions is the ability to work with a variety of transport technologies. The WX Framework and the WX and WXC product lines provide flexible deployment options to enable interoperability with any number of WAN architectures. The WX and WXC platforms also include features to ensure interoperability with transports that have specific traffic marking requirements. Multi-Protocol Label Switching (MPLS) is one transport that requires specific preservation of packet marking.

## Overview of WAN Services Transports

Historically, corporate WAN data services have run over a number of different circuit-oriented network transports, including point-to-point leased lines, X.25, ATM, frame relay, DSL, and IP VPNs, including MPLS VPNs.

Point-to-point leased lines used dedicated physical circuits, leased by the telco or PTT to the customer. Customer-owned routers reside at each end of the leased line, and a Layer 2 encapsulation such as PPP or HDLC is used to transport TCP/IP or other higher-layer traffic. These leased lines have fixed, non-shared bandwidth capacity and are point-to-point connections.

X.25 networks introduced the notion of “virtual circuits,” allowing multiple virtual connections to traverse a single physical circuit and creating more flexibility for site interconnection. These virtual circuits are available as either “permanent virtual circuits” (PVCs), which map two specific sites together, or “switched virtual circuits” (SVCs), which have a call-setup facility that allows them to dynamically set up connections to many different sites, somewhat analogous to phone calls on the telephone network. The virtual circuits are still fixed in their bandwidth capacity.

ATM networks use a model similar to X.25’s, but they add mechanisms to provide strict guarantees for quality of service and better integration for voice and video. These circuits also provide fixed bandwidth capacity.

The most common transport is frame relay. Like X.25, frame relay is also based on virtual circuit networks, but it has the added feature of providing variable bandwidth. Typically customers pay for a committed rate and are allowed to burst to a higher rate if the network is not congested. This architecture introduced the notion of per-circuit, variable WAN bandwidth. Frame relay networks are the dominant WAN network transport in use today across the world, but their growth is tapering off as new network transport options emerge. Another circuit-oriented, variable-rate network transport that has been growing steadily in the last few years is DSL.

Most recently, IP VPNs have become available in many regions across the world as an alternative to the dedicated circuit model. With IP VPNs, either IP tunnels or label-switched paths can be used to form multiple logical connections over a single physical connection. This architecture simplifies circuit management and greatly increases flexibility of IP connectivity. In an IP VPN, the customer still uses a circuit-oriented network transport to access the network, but the customer needs only a single circuit to access locations throughout the entire network.

Both connection-oriented and connectionless options are available over IP VPNs. The connection-oriented approach uses an IP-based tunneling protocol such as IPsec, PPTP, or

L2TP to create multiple, point-to-point connections across the single physical or virtual circuit used to access the network. While this approach provides for any-to-any connectivity, configuration is a challenge. A full-mesh configuration is needed, which requires complete reconfiguration of the VPN when any new sites come online. Because of this configuration challenge, IP-based tunneling VPNs are typically modeled as hub-and-spoke to isolate the number of locations where configuration changes are needed.

The connectionless approach uses a sub-IP layer label-switching protocol, MPLS, to create multiple logical connections across the single physical or virtual circuit used to access the network. With MPLS, point-to-point connections are not needed – instead, a dynamic label-switching mechanism provides connectionless connectivity. This design allows MPLS networks to provide any-to-any connectivity without having to reconfigure VPNs when new sites come online.

IP VPNs are gaining in popularity, and the easier-to-configure connectionless ones, based on MPLS, are especially popular. A number of the world's leading telcos offer MPLS VPNs, including AT&T, British Telecom, Deutsche Telecom, and many others.

While MPLS VPNs provide a number of configuration advantages, they also require specific packet markings to operate correctly. Many WAN optimization platforms cannot support these required markings. Juniper Networks understands the importance of successfully interoperating with telco operators' MPLS networks and has taken care to support the necessary features in its WAN optimization and application acceleration platforms.

## MPLS Integration Issues

At the highest level, MPLS networks are transparent to LAN devices on the corporate networks they serve, including Juniper WX and WXC devices. At the edge of the corporate network, a WAN router is configured to speak MPLS on its WAN-facing interface(s). This WAN router is often referred to as an MPLS "Provider Edge" router (MPLS PE). On this MPLS PE router, an interior routing protocol and label-distribution mechanism dynamically bind information about IP reachability with that of available label-switched paths through the MPLS network. In a sense, the label-switched path provides a transport service somewhat similar to a switched virtual circuit. On the LAN-facing interface(s) of the MPLS PE, only an interior routing protocol needs to run (such as OSPF or EIGRP).

Although the control- and forwarding-plane operations of the MPLS network are entirely transparent to LAN devices behind the MPLS PE router, as the WX and WXC platforms are, many MPLS providers offer class of service guarantees for different classes of traffic that traverse the WAN. The challenge then is to classify the different traffic classes appropriately either at or before the MPLS PE router. Several approaches are available for this classification.

It's possible to map applications directly to classes at the packet level by using application recognition capabilities of the router and marking the packet appropriately. The router can also mark traffic based on incoming interface or IP subnet. Marking techniques include setting the IP Precedence field of the IP ToS byte in the IP header or setting the DiffServ Code Point (DSCP) in the IP header.

A WX or WXC device compresses the raw IP flows destined for the WAN and encapsulates them in UDP or IPcomp tunnel packets. In cases where the MPLS network does not use class of service guarantees, the WX and WXC platforms do not need to make any special accommodations. But to integrate successfully with an MPLS network that does use class of service guarantees, the packet modification performed by the WX or WXC platform must take into account some requirements of that MPLS network.

## Tunnel Class of Service Marking

If the MPLS PE router is configured to use IP ToS or DSCP values in the IP header to determine which MPLS class of service to assign, then a device somewhere in the network must be responsible for classifying the traffic properly and applying the proper IP ToS or DSCP value to each packet. All WX and WXC devices can provide this class of service marking feature.

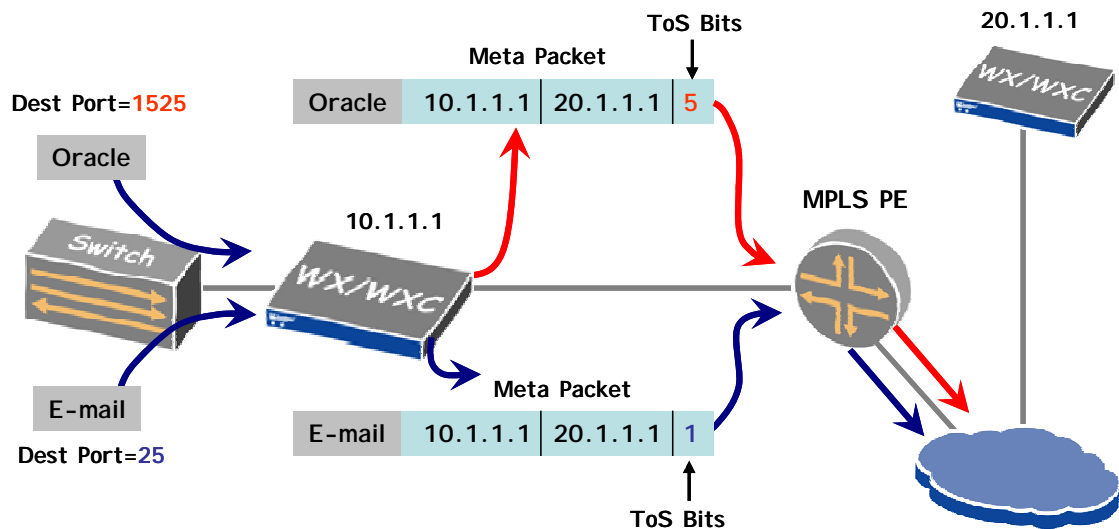



Figure 1: Tunnel Class of Service Marking

The WX/WXC platform can be configured to group applications into traffic classes and then assign the proper IP ToS or DSCP value to each class of traffic (see Figure 1). The WX or WXC platform will then place traffic from the same classes into the same tunnel packets, and the tunnel packet's IP header contains the proper IP ToS or DSCP value. The result is that the MPLS PE simply looks at the tunnel packet's IP header to determine the proper class of service for a given class of traffic. This approach is completely transparent to the MPLS PE router and the MPLS network, and it leverages the easy-to-use WX software's GUI interface to map applications to classes.

## Tunnel Class of Service Preservation

In some cases, the class of service marking required by the MPLS PE router may have happened before the packet arrives at the WX or WXC device. In this case, the WX or WXC platform must retain the IP ToS or DSCP marking on the original IP packets. The platform copies the IP ToS or DSCP value from the IP header of the original packet to the tunnel packet's IP header (see Figure 2). All packets that share the same IP ToS or DSCP value travel in the same tunnels. This operation is simple and automatic for the WX/WXC device – IT can enable it by clicking a single checkbox in the GUI.



**Figure 2: Tunnel Class Preservation**

Note that it is also possible to preserve the original IP ToS or DSCP marking of the original IP packet when it is received by the WX or WXC platform on the far side of the MPLS network. This option, too, requires simply clicking a single checkbox in the GUI.

## Tunnel Application Marking

If the MPLS PE router is configured to recognize applications and map them to the proper MPLS class of service, then the WX or WXC device cannot obscure the application port numbers contained in the IP header. The router needs to assign applications to a given traffic class, but routers cannot see tunneled applications. With tunneled applications, the application port information is obscured because the original packet is inside the tunnel. Overcoming this problem requires a technique called Tunnel Application Marking, or Application Visibility mode.

In Application Visibility mode (see Figure 3), the WX or WXC device compresses the original IP packets and places them in UDP tunnels. The platform then creates different UDP tunnels for each application and it copies the TCP or UDP port number in the original IP packet to the tunnel packet's IP header. When the MPLS PE receives these tunnel packets, it can still recognize the applications based on the tunnel packet port numbers.

Note that if the application is a TCP application, the MPLS PE will need a simple configuration modification to look for the UDP equivalent of the TCP port. For example, if a router's access list has been defined to identify an application based on TCP port 80, it will need one additional rule entered to also recognize port 80 in UDP packets.

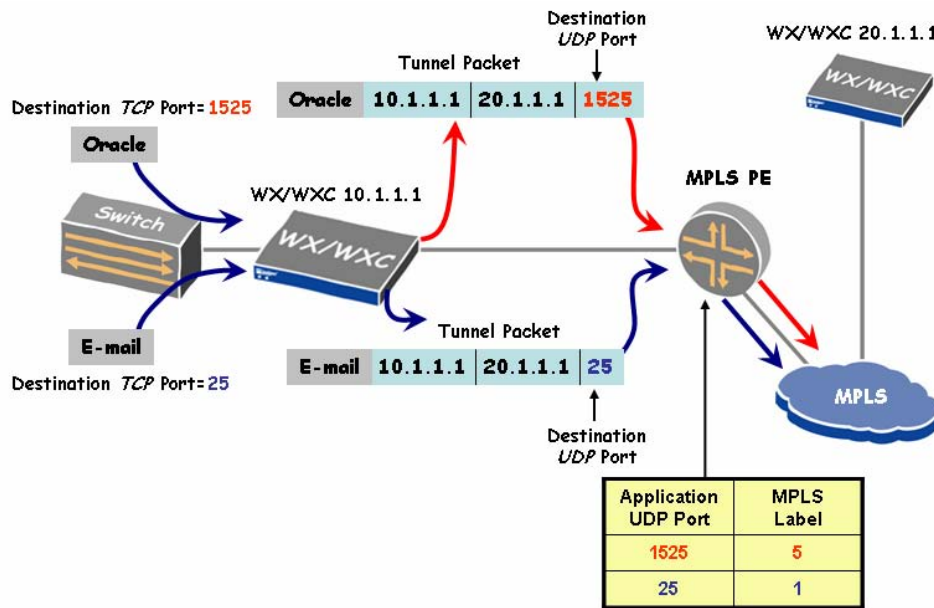


Figure 3: Application Visibility Mode

## Summary

MPLS networks provide a very flexible WAN transport service that is largely transparent to the LANs they serve. For WAN optimization platforms to completely integrate with MPLS networks, they need to provide a wide range of

options for mapping traffic to the appropriate MPLS class of service. Juniper Networks devices offer a rich set of options for these and other network integration issues, enabling IT to deploy the WX and WXC devices transparently in any MPLS environment.

### Sample Router Configuration Change

*Before Tunnel Application Marking:*

```
access-list 101 permit tcp any 10.1.1.0.0.255 eq 80
```

*After Tunnel Application Marking:*

```
access-list 101 permit tcp any 10.1.1.0.0.255 eq 80
access-list 101 permit udp anv 10.1.1.0.0.255 eq 80
```