

## Chapter 21

# Configuring CoS for MPLS

When IP traffic enters a label-switched path (LSP) tunnel, the ingress router marks all packets with a class-of-service (CoS) value, which is used to place the traffic into a transmission priority queue. On the router, each interface has up to eight transmit queues. The CoS value is encoded as part of the Multiprotocol Label Switching (MPLS) header and remains in the packets until the MPLS header is removed when the packets exit from the egress router. The routers within the LSP utilize the CoS value set at the ingress router. The CoS value is encoded by means of the CoS bits (also known as the EXP or experimental bits).

MPLS class of service works in conjunction with the router's general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin, and to configure congestion avoidance using random early detection (RED).

The next-hop label-switched router (LSR) uses the default classification shown in Table 37.

**Table 37: LSR Default Classification**

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low
111	network-control	high

To configure CoS for MPLS, you can include the following statements at the [edit class-of-service] hierarchy level of the configuration:

If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

To specify a CoS value for packets in an LSP, include the `class-of-service` statement:

```
class-of-service cos-value;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit protocols mpls interface *interface-name* label-map *label-value*]
- [edit protocols mpls label-switched-path *path-name*]
- [edit protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit protocols mpls static-path *prefix*]
- [edit protocols rsvp interface *interface-name* link-protection]
- [edit protocols rsvp interface *interface-name* link-protection bypass *destination*]
- [edit logical-routers *logical-router-name* protocols mpls]
- [edit logical-routers *logical-router-name* protocols mpls label-switched-path *path-name*]
- [edit logical-routers *logical-router-name* protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit logical-routers *logical-router-name* protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit logical-routers *logical-router-name* protocols mpls static-path *prefix*]
- [edit logical-routers *logical-router-name* protocols mpls interface *interface-name* label-map *label-value*]
- [edit logical-routers *logical-router-name* protocols rsvp interface *interface-name* link-protection]
- [edit logical-routers *logical-router-name* protocols rsvp interface *interface-name* link-protection bypass *destination*]

The `class-of-service` statement at the [edit protocols mpls label-switch-path] hierarchy level assigns an initial EXP value for the MPLS shim header of packets in the LSP. This value is initialized at the ingress router only and overrides the rewrite configuration established for that forwarding class. However, the CoS processing (weighted round robin [WRR] and RED) of packets entering the ingress router is not changed by the `class-of-service` statement on an MPLS LSP. Classification is still based on the behavior aggregate (BA) classifier at the [edit class-of-service] hierarchy level or the multifield (MF) classifier at the [edit firewall] hierarchy level.

We recommend configuring all routers along the LSP to have the same input classifier for EXP, and, if a rewrite rule is configured, all routers should have the same rewrite configuration. Otherwise, traffic at the next LSR might be classified into a different forwarding class, resulting in a different EXP value being written to the EXP header.

For more information, see the *JUNOS MPLS Applications Configuration Guide*.