Chapter 13
LDP Overview

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an end point at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or they might have an end point at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

This chapter discusses the following topics:

Overview

LDP associates a set of destinations (route prefixes and router addresses) with each data link LSP. This set of destinations is called the Forwarding Equivalence Class (FEC). These destinations all share a common data LSP path egress and a common unicast routing path. Each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This forms a tree of LSPs that converge on the egress router.

You can implement Virtual Private Networks (VPNs) using MPLS for tunneling. This allows the use of overlapping address spaces by different VPNs. Some of these MPLS-based approaches to VPNs support only LDP for signaling. With JUNOS implementation of LDP, and Juniper Networks routers at the core of a network, you can implement edge devices that support VPNs using LDP signaling for MPLS.

LDP Standards

LDP is described in: *Label Distribution Pr otocol (LDP)—Version 1 Functional Specification*, Internet draft draft-ietf-mpls-ldp-06.txt.

To access Internet RFCs and drafts, go to the IETF web site www.ietf.org.

JUNOS LDP Protocol Implementation

The JUNOS implementation of LDP supports LDP Version 1. The JUNOS software supports a simple mechanism for tunneling between routers in an IGP, to eliminate the required distribution of external routes within the core. JUNOS software version 4.1 allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP, but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary.

LDP Operation

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. To do this, include the family mpls statement at the [edit interfaces *interface* unit *unit*] hierarchy level.

```
[edit interfaces interface unit unit]
family mpls;
```

LDP Label Filtering

You can apply policy filters to labels received from other routers via LDP. This provides a mechanism to administratively limit the establishment of LSPs. To do this, include the import statement at the [edit protocols ldp] hierarchy level.

```
[edit protocols ldp]
import [ policy-names ];
```

Tunneling LDP LSPs in RSVP LSPs

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP will be tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP will automatically establish sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than being carried through the LSP. This allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

Label Operations

Figure 17 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see “Label Description” on page 17.) The shaded inner oval represents the RSVP domain, while the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 17: Swap and Push Label Operation when Tunneling LDP LSPs in RSVP LSPs

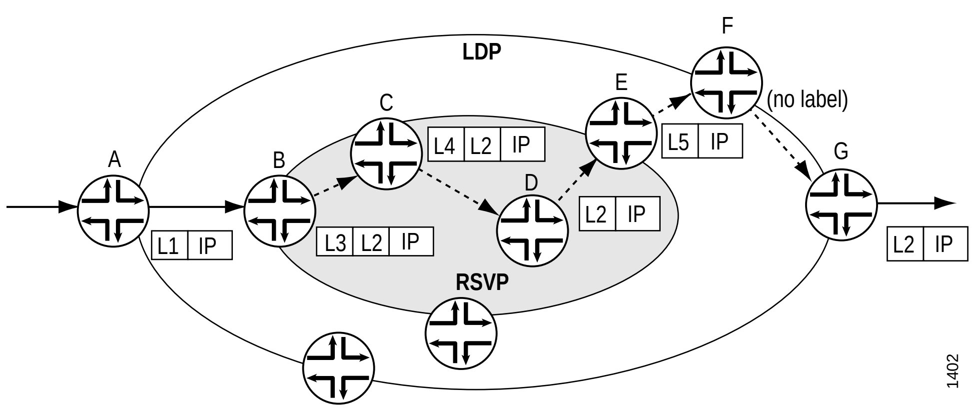

Figure 18 depicts double push label operation (L1L2), which is used when the ingress router (A) of the LDP and the RSVP are the same router. Note that router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by router D.

Figure 18: Double Push Label Operation when Tunneling LDP LSPs in RSVP LSPs

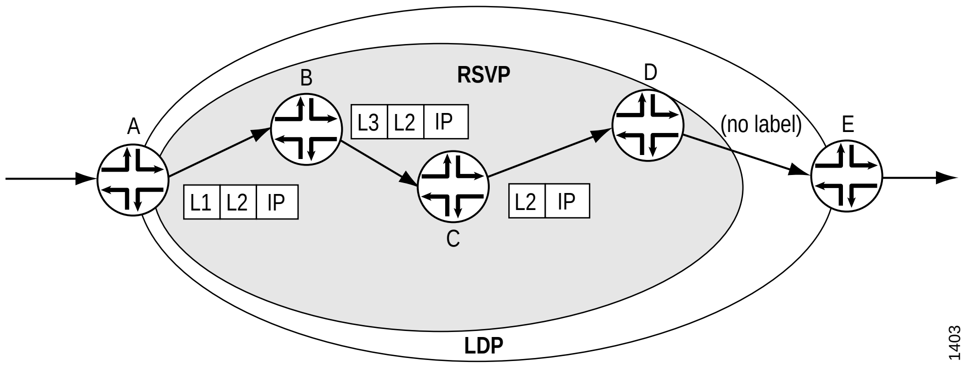

Restrictions for LDP over RSVP

The IGP shortcut computation imposes some restrictions on the network topology allowed. All the routers in the traffic engineered core and in the surrounding LDP cloud must belong to the same OSPF area or IS-IS level. Using multiple areas or levels prevents the IGP shortcut computation from finding an RSVP LSP next hop. As a result, you cannot use a label from a remote LDP session for this router.

If all the routers do not belong to the same area or level, then traffic engineering shortcuts must be explicitly enabled in the IGP. (For more information, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.) You cannot use RIP as an IGP for shortcuts.

LDP Message Types

LDP uses several types of messages to establish and remove mappings, and to report errors. All LDP messages have a common structure that uses a type-length-value (TLV) encoding scheme.

Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending the hello message periodically. This hello message is transmitted as a UDP packet to the LDP port at the group multicast address for all routers on the subnet.

Session Messages

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure completes successfully, the two routers are LDP peers, and can exchange advertisement messages.

Advertisement Messages

Advertisement messages create, change, and delete label mappings for Fast Ethernet Channels (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

Notification Messages

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

Error notifications signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.

Advisory notifications pass a router information about the LDP session or the status of some previous message received from the peer.