Foundry NetIron Service Provider Configuration and Management Guide

FOUNDRY NETWORKS

2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
Tel 408.586.1700
Fax 408.586.1900

November 2002

Contents

Chapter 1
Getting Started

Introduction

This guide describes the Layer 2 Switch, Layer 3 Switch, and ServerIron product families and features from Foundry Networks. Procedures are provided for installing the hardware and configuring the software. The software procedures show how to perform tasks using the Command Line Interface (CLI).

This guide also describes how to monitor Foundry products using statistics and summary screens.

Audience

This manual is designed for system administrators with a working knowledge of Layer 2 and Layer 3 switching and routing.

If you are using a Foundry Layer 3 Switch, you should be familiar with the following protocols if applicable to your network – IP, RIP, OSPF, IS-IS, BGP4, MBGP, MPLS, IGMP, PIM, DVMRP, IPX, AppleTalk, FSRP, VRRP, and VRRPE.

Nomenclature

This guide uses the following typographical conventions to show information:

Italic	highlights the title of another publication and occasionally emphasizes a word or phrase.
Bold	highlights a CLI command.
Bold Italic	highlights a term that is being defined.

NOTE: A note emphasizes an important fact or calls your attention to a dependency.

WARNING: A warning calls your attention to a possible hazard that can cause injury or death.

CAUTION: A caution calls your attention to a possible hazard that can damage equipment.

Related Publications

The following Foundry Networks documents supplement the information in this guide.

- *Foundry Switch and Router Installation and Basic Configuration Guide* – provides hardware and software installation information, and configuration information for system-level features.
- *Foundry Security Guide* – provides procedures for securing management access to Foundry devices and for protecting against Denial of Service (DoS) attacks.
- *Foundry Enterprise Configuration and Management Guide* – provides configuration information for enterprise routing protocols including IP, RIP, IP multicast, OSPF, BGP4, VRRP and VRRPE.
- *Foundry Switch and Router Command Line Interface Reference* – provides a list and syntax information for all the Layer 2 Switch and Layer 3 Switch CLI commands.
- *Foundry Diagnostic Guide* – provides descriptions of diagnostic commands that can help you diagnose and solve issues on Layer 2 Switches and Layer 3 Switches.

To order additional copies of these manuals, do one of the following:

- Call 1.877.TURBOCALL (887.2622) in the United States or 1.408.586.1881 outside the United States.
- Send email to info@foundrynet.com.

What’s New In This Edition?

This edition describes the following software release:

- 07.6.01

This release applies to the following products:

- NetIron Internet Backbone router
- BigIron with M2 (Management II) or higher modules
- BigIron with Velocity Management Module version I (VM1)
- FastIron II, FastIron II Plus, and FastIron III with M2 or higher modules
- FastIron 4802

For a list of the enhancements, see the “Getting Started” chapter in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

NOTE: If you want documentation specifically for a 07.1.x release, see the January, 2001 edition of the manuals and the release notes for the release you are using. For the 07.2.06 release or a 07.3.x release, see the June, 2001 edition and the release notes.

How to Get Help

Foundry Networks technical support will ensure that the fast and easy access that you have come to expect from your Foundry Networks products will be maintained.

Web Access

- http://www.foundrynetworks.com

Email Access

Technical requests can also be sent to the following email address:

- support@foundrynet.com

Telephone Access

- 1.877.TURBOCALL (887.2622)   United States

• 1.408.586.1881	Outside the United States

Warranty Coverage

Contact Foundry Networks using any of the methods listed above for information about the standard and extended warranties.

Chapter 2
NetIron IP-Only Software

The NetIron Internet Backbone router is optimized for IP-only environments. Each NetIron Internet Backbone router shipped from the factory comes with the following image files and with an empty configuration (no startup-config file).

- B2Pxxxxx.bin
- N2Pxxxxx.bin

The installation and configuration procedures for the NetIron Internet Backbone router are the same as those for the BigIron Layer 3 Switch. However, depending on the ISP-only image you run on the device, some of the system defaults are different.

System Defaults

Most of the system defaults in these IP-only images are the same as in the enterprise image (B2Rxxxxx.bin). However, the following parameters have different default values in the B2P and N2P images than they do in the B2R image:

- Port state – When you boot using the N2P image, all ports are disabled by default. All ports are enabled by default if you boot using B2P.
- Layer 2 support – When you boot using B2P or N2P, Layer 2 support is disabled. The device is an IP router only and does not perform Layer 2 switching. Layer 2 support is enabled by default if you boot using the B2R image.

The defaults described above apply to new NetIron Internet Backbone routers that you boot without a startup-config file. However, if you are upgrading a device that has a startup-config file, the default for Layer 2 forwarding depends on the image you are booting:

- B2P – Layer 2 support is enabled. This is the state if the startup-config file does not contain a **route-only** command to disable Layer 2 support.
- N2P – Layer 2 support is disabled. This is the state if the startup-config file does not contain a **no route-only** command to enable Layer 2 support.

Determining the State of Layer 2 Support

To determine the state of Layer 2 support, enter the following command at any level of the CLI:

```
BigIron# show ip
Global Settings
  ttl: 64, arp-age: 10, bootp-relay-max-hops: 4
  router-id : 192.168.1.11
  enabled : UDP-Broadcast-Forwarding  Source-Route  Load-Sharing  RARP  RIP
  disabled: Route-Only  Directed-Broadcast-Forwarding  BGP4 IRDP  Proxy-ARP  RIP
-Redist  OSPF  DVMRP  FSRP  VRRP  VRRP-Extended
```

If “Route-Only” is listed following “enabled”, Layer 2 support is disabled. If “Route-Only” is listed following “disabled”, Layer 2 support is enabled.

You also can display the running-config to see whether it contains a **route-only** or **no route-only** command.

- If you display the running-config on a device that you booted without a startup-config file:
  - B2P – The running-config contains the **route-only** command, indicating that Layer 2 support is disabled.
  - N2P – The running-config does not contain the **route-only** or **no route-only** command. However, Layer 2 support is disabled.
- If you display the running-config on a device that you booted using a startup-config file:
  - B2P – If Layer 2 support is disabled, the **route-only** command is displayed. If Layer 2 support is enabled, the **no route-only** command is displayed.
  - N2P – If Layer 2 support is enabled, the **no route-only** command is displayed. If Layer 2 support is disabled, the **route-only** command is displayed.

NOTE: This does not apply to the B2R image. The running-config contains the **route-only** command, if Layer 2 support is disabled. If Layer 2 support is enabled (the default), the running-config does not contain the **route-only** or the **no route-only** command.

IP-Only Features

The IP-only image files on the NetIron Internet Backbone router support all the IP features supported in the enterprise image file. In addition, the IP-only images support the following features:

- ***Multiprotocol Label Switching (MPLS)*** can be used to direct packets through a network over a predetermined path of routers. Forwarding decisions in MPLS are based on the contents of a label applied to the packet, instead of information in the packet’s IP header. See “Configuring MPLS” on page 3-1.
- ***Traffic engineering*** is the ability to direct packets through a network efficiently, using information gathered about network resources. When used as an application of MPLS, traffic engineering involves creating paths that make the best use of available network resources, avoiding points of congestion and making efficient use of high bandwidth interfaces. Packets travelling over these paths are forwarded using MPLS. See “Configuring MPLS” on page 3-1.
- ***MPLS Virtual Leased Line (VLL)*** is a method for providing point-to-point Ethernet/VLAN connectivity over an MPLS domain. This functionality is outlined in the IETF draft-martini documents.
- ***Intermediate System to Intermediate System (IS-IS)*** – A link-state Interior Gateway Protocol (IGP) that routers (intermediate systems) can use to exchange routes within a single routing domain. IS-IS is based on the International Standard for Organization/International Electrotechnical Commission (ISO/IEC) Open Systems Interconnect (OSI) networking model, and describes communication within the Networking layer of the model. See “Configuring IS-IS” on page 6-1.

Chapter 3
Configuring MPLS

This chapter explains how to configure ***Multiprotocol Label Switching (MPLS)*** on a Foundry device for traffic engineering purposes. MPLS can be used to direct packets through a network over a predetermined path of routers. Forwarding decisions in MPLS are based on the contents of a label applied to the packet, instead of information in the packet's IP header.

Traffic engineering is the ability to direct packets through a network efficiently, using information gathered about network resources. When used as an application of MPLS, traffic engineering involves creating paths that make the best use of available network resources, avoiding points of congestion and making efficient use of high bandwidth interfaces. Packets travelling over these paths are forwarded using MPLS.

Overview

This chapter is divided into the following sections:

- "MPLS Hardware and Software Support on Foundry Devices" on page 3-2 lists the hardware and software necessary to run MPLS on Foundry devices. This section also lists the IETF RFCs and Internet Drafts supported by Foundry's implementation of MPLS.
- "How MPLS Works" on page 3-3 explains basic concepts about MPLS, including how packets are forwarded along a Label Switched Path (LSP), the difference between static and signalled LSPs, and how an MPLS label header is encoded.
- "Using MPLS in Traffic Engineering" on page 3-8 explains the process that takes place to create and activate traffic-engineered LSPs.
- "Using Traffic Engineered LSPs Between IBGP Neighbors" on page 3-17 shows how traffic-engineered LSPs can be used as "shortcuts" between IBGP neighbor routers in an Autonomous System (AS).
- "Using Traffic Engineered LSPs Within an AS" on page 3-21 explains how traffic-engineered LSPs can be used as shortcuts to IGP destinations within an AS.
- "Configuring MPLS" on page 3-23 describes how to set up MPLS on Foundry devices using the Command Line Interface (CLI).
- "Displaying MPLS and RSVP Information" on page 3-41 describes the commands used to display information about an MPLS configuration.
- "MPLS Sample Configurations" on page 3-58 shows diagrams of typical MPLS configurations and the CLI commands used for implementing them.

MPLS Hardware and Software Support on Foundry Devices

This section details the combination of hardware chassis, management module, MPLS-capable modules, and software image required to run Foundry's implementation of MPLS.

Hardware Chassis

MPLS is supported on the NetIron 400, NetIron 800, and NetIron 1500 Chassis devices.

Management Module

Foundry's implementation of MPLS requires the Management IV module or the VM1 module.

The Management IV module must be running boot code version 07.5.00 (M2B07500.bin) or later.

MPLS-Capable Modules

MPLS-capable modules are defined as modules that can perform the following tasks:

- Transmit and receive RSVP and LDP messages for the purpose of establishing LSPs
- Transmit and receive data packets with MPLS label encapsulation
- Convert between labeled and unlabeled packet formats.

The MPLS-capable modules on a NetIron constitute the portion of the router that can belong to an MPLS domain. From the perspective of a Label Edge Router (LER), the MPLS-capable modules are those that aggregate unlabeled packets coming from non-MPLS-capable modules and forward them into LSP tunnels and vice versa. A NetIron chassis can contain any combination of MPLS-capable and non-MPLS-capable modules.

Note that any port on an MPLS-capable module can perform both MPLS forwarding and regular Layer 3 IP forwarding simultaneously. An MPLS-capable module can also perform LER functions for some packet flows and LSR (Label Switching Router) functions for others.

A port on an MPLS-capable module begins performing MPLS functions once you explicitly configure MPLS on the router and enable the port as an MPLS interface.

MPLS Support with Management IV Modules

When a NetIron Chassis device is using a Management IV module, the following Foundry modules can function as MPLS-capable modules:

- OC-3 POS
- OC-12 POS
- OC-48 POS (NPA and non-NPA)

When a NetIron Chassis device is using a Management IV module, the following Foundry modules can function as non-MPLS-capable, regular Layer 3 modules:

- 10/100 Ethernet
- Gigabit Ethernet
- OC-3 POS
- OC-12 POS
- OC-48 POS (NPA and non-NPA)

MPLS Support with VM1 Modules

Starting with release 07.6.01, MPLS is supported on Ethernet interfaces on NetIron 1500, NetIron 800, and NetIron 1500 Chassis devices with VM1 modules installed.

For NetIron Chassis devices with VM1 modules, the same MPLS features that were available in pre-07.6.01 releases are now available on Ethernet interfaces, including RSVP signalled LSPs, static LSPs, BGP/IGP shortcuts, and draft-Martini VLLs. The new MPLS features in release 07.6.01 are also supported on Ethernet interfaces on devices with VM1 modules.

When a NetIron Chassis device is using a VM1 module, the following Foundry modules can function as MPLS-capable modules:

- 10/100 Ethernet
- Gigabit Ethernet
- OC-3 POS
- OC-12 POS
- OC-48 POS (NPA and non-NPA)

When a NetIron Chassis device is using a VM1 module, the following Foundry modules can function as non-MPLS-capable, regular Layer 3 modules:

- 10/100 Ethernet
- Gigabit Ethernet
- OC-3 POS
- OC-12 POS
- OC-48 POS (NPA and non-NPA)

Software Support

The N2M image (N2M*****.bin), contains MPLS. Since the N2M image is too large to fit in flash memory, you must boot the N2M image from a PCMCIA flash card.

New NetIron Internet Backbone routers are shipped from the factory with the N2P and B2P images in flash memory and with the N2M image on the PCMCIA flash card. See “Note to MPLS Users” in the *IronWare Release Notes* for information on booting the N2M image from a PCMCIA flash card.

IETF RFC and Internet Draft Support

Foundry’s implementation of MPLS supports the following IETF RFCs and Internet Drafts.

MPLS

RFC 2702 Requirements for Traffic Engineering Over MPLS

RFC 3031 Multiprotocol Label Switching Architecture

RFC 3032 MPLS Label Stack Encoding

draft-ietf-mpls-icmp-02.txt ICMP Extensions for MultiProtocol Label Switching

RSVP

RFC 2205 Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification

draft-ietf-mpls-rsvp-lsp-tunnel-08.txt RSVP-TE: Extensions to RSVP for LSP Tunnels

OSPF-TE

RFC 2370 The OSPF Opaque LSA Option

draft-katz-yeung-ospf-traffic-03.txt Traffic Engineering Extensions to OSPF

LDP for Tunnel LSPs (New in Release 07.6.01)

RFC 3036 – LDP Specification

How MPLS Works

MPLS uses a ***label switching*** forwarding method to direct packets through a network. In label switching, a packet is assigned a label and passes along a predetermined path of routers. Forwarding decisions are based on the contents of the label, rather than information in the packet’s IP header.

The following sections describe these basic MPLS concepts:

- How packets are forwarded through an MPLS domain
- The kinds of Label Switched Paths (LSPs) that can be configured on a Foundry device
- The components of an MPLS label header

How Packets Are Forwarded Through an MPLS Domain

An ***MPLS domain*** consists of a group of MPLS-enabled routers, called ***LSRs*** (Label Switching Routers). In an MPLS domain, packets are forwarded from one MPLS-enabled router to another along a predetermined path, called an ***LSP*** (Label Switched Path). LSPs are one-way paths between MPLS-enabled routers on a network. To provide two-way traffic, you configure LSPs in each direction.

The LSRs at the headend and tailend of an LSP are known as ***LERs*** (Label Edge Routers). The LER at the headend, where packets enter the LSP, is known as the ***ingress LER***. The LER at the tailend, where packets exit the LSP, is known as the ***egress LER***. Each LSP has one ingress LER and one egress LER. Packets in an LSP flow in one direction: from the ingress LER towards the egress LER. In between the ingress and egress LERs there may be one or more ***transit LSRs***. A Foundry device enabled for MPLS can perform the role of ingress LER, transit LSR, or egress LER in an LSP. Further, a Foundry device can serve simultaneously as an ingress LER for one LSP, transit LSR for another LSP, and egress LER another LSP.

Figure 3.1 depicts an MPLS domain with a single LSP consisting of three LSRs: an ingress LER, a transit LSR, and an egress LER.

Figure 3.1 Label switching in an MPLS domain

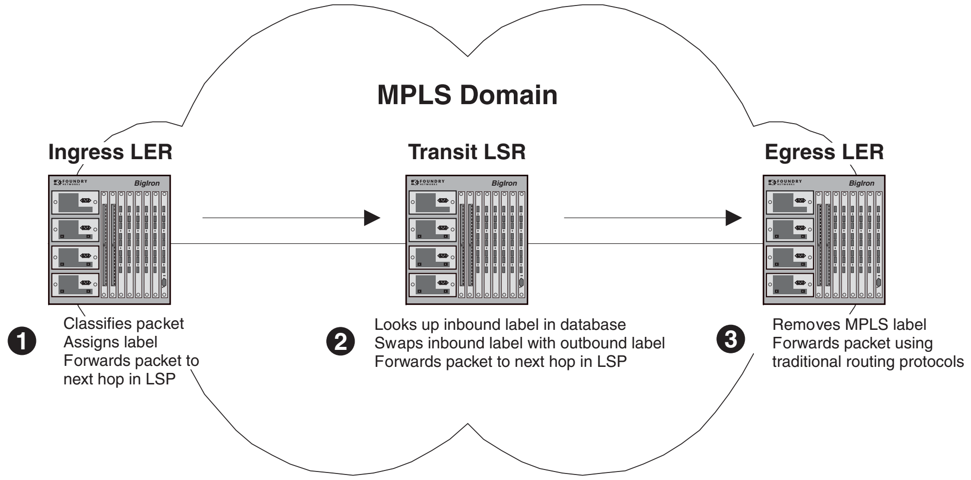

Label switching in an MPLS domain works as follows:

1. The Ingress LER receives a packet and pushes a label onto it.

   When a packet is received on an MPLS-enabled interface, the Foundry device determines to which LSP (if any) the packet should be assigned. Specifically, the device determines to which Forwarding Equivalence Class (FEC) the packet belongs. A FEC is simply a group of packets that should all be forwarded in the same way. For example, a FEC could be defined as all packets whose destination address matches a specified IP address prefix. FECs are mapped to LSPs. If a packet belongs to a FEC, and an LSP is mapped to that FEC, the packet is assigned to the LSP.

   When a packet is assigned to an LSP, the Foundry device, acting as an ingress LER, applies (pushes) a ***label*** onto the packet. A label is a 32-bit, fixed-length identifier that is significant only to MPLS. See “MPLS Label Header Encoding” on page 3-6 for specific information about the contents of a label. From this point until the packet reaches the egress LER at the end of the path, the packet is forwarded using information in its label, not information in its IP header. The packet’s IP header is not examined again as long as the packet traverses the LSP.

On the ingress LER, the label is associated with an outbound interface. Once assigned a label, the packet is forwarded over the outbound interface to the next router in the LSP.

2. A transit LSR receives the labelled packet, swaps the label, and forwards the packet to the next LSR.

   In an LSP, there can be zero or more transit LSRs between the ingress and egress LERs. A transit LSR swaps labels on an MPLS packet and forwards the packet to the next router in the LSP.

   When a transit LSR receives an MPLS packet, it looks up the label in its ***MPLS forwarding table***. This table maps the label and inbound interface to a new label and outbound interface. The transit LSR replaces the old label with the new label and sends the packet out the outbound interface specified in the table. This process repeats at each transit LSR until the packet reaches either the egress LER (for static LSPs), or the next-to-last LSR in the LSP (for signalled LSPs).

   Figure 3.2 shows an example of the label swapping process on a transit LSR.

Figure 3.2 Label swapping on a transit LSR

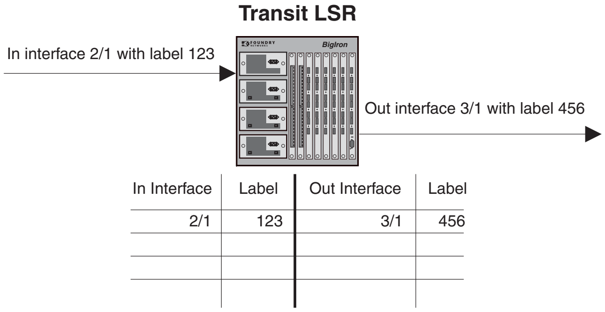
Transit LSR
In interface 2/1 with label 123
Out interface 3/1 with label 456
BigIron

In Interface	Label	Out Interface	Label
2/1	123	3/1	456

   In this example, a packet comes into interface 2/1 with label 123. The transit LSR then looks up this interface-label pair in its MPLS forwarding table. The inbound interface-label pair maps to an outbound-interface-label pair – in this example, interface 3/1 with label 456. The LSR swaps label 123 with label 456 and forwards the packet out interface 3/1.

3. The egress LER receives labelled packet, pops label, and forwards IP packet.

   When the packet reaches the egress LER, the MPLS label is removed (called ***popping*** the label), and the packet can then be forwarded to its destination using standard hop-by-hop routing protocols. On signalled LSPs, the label is popped at the penultimate (next to last) LSR, rather than the egress LER. See “Penultimate Hop Popping” on page 3-6 for more information.

Types of LSPs

An LSP in an MPLS domain can be either ***static*** or ***signalled***.

Static LSPs

Static LSPs are configured manually on each LSR in the LSP. No signalling protocol is used. To establish a static LSP, you configure the ingress LER, transit LSRs, and egress LER, manually specifying the labels to be applied at each hop.

See “Setting up Static LSPs” on page 3-35 for more information.

Signalled LSPs

Signalled LSPs are configured only at the ingress LER. When the LSP is enabled, RSVP signalling messages travel to each LSR in the LSP, reserving resources and causing labels to be dynamically associated with interfaces. When a packet is assigned to a signalled LSP, it follows a pre-established path from the LSP’s ingress LER to its egress LER. This path can be one of the following:

- A path that traverses an explicitly specified set of MPLS routers
- The IGP shortest path across the MPLS domain, determined from local routing tables

- A traffic-engineered path calculated by the Foundry device using constraints such as bandwidth reservations, administrative groups, and network topology information

For more information, see “How CSPF Calculates a Traffic-Engineered Path” on page 3-10, “How RSVP Establishes a Signalled LSP” on page 3-11, and “Setting Up Signalled LSPs” on page 3-27.

Penultimate Hop Popping

On signalled LSPs, the MPLS label is popped at the next-to-last LSR in the LSP, instead of the egress LER. This action is called ***penultimate hop popping***. Penultimate hop popping improves forwarding efficiency by allowing the egress LER to avoid performing both a MPLS forwarding table lookup and an IP forwarding table lookup for each packet exiting the LSP. Instead, the MPLS label is popped at the penultimate (next-to-last) LSR, and the packet is forwarded to the egress LER with no MPLS encoding. The egress LER, in fact, does not recognize the packet as emerging from an LSP.

Figure 3.3 illustrates the operation that takes place at the penultimate LSR in an LSP.

Figure 3.3 Penultimate hop popping

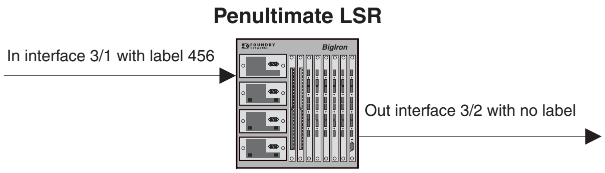
Penultimate LSR
In interface 3/1 with label 456
FOUNDRY NETWORKS
BigIron
Out interface 3/2 with no label

In Interface	Label	Out Interface	Label
3/1	456	3/2	[none]

When an LSR receives an MPLS packet, it looks up the label in its MPLS forwarding table. Normally, this table maps the label and inbound interface to a new label and outbound interface. However, when this is the penultimate LSR in an LSP, the label and inbound interface map only to an outbound interface. The penultimate LSR pops the label and forwards the packet – now a regular IP packet – out the outbound interface. When the packet reaches the egress LER, there is no indication that it had been forwarded over an LSP. The packet is forwarded using standard hop-by-hop routing protocols.

NOTE: Penultimate hop popping is always performed on signalled LSPs, but never performed on static LSPs.

MPLS Label Header Encoding

The following diagram illustrates the structure of the 32-bit MPLS label header. When a packet enters an LSP, the ingress LER pushes a label onto the packet.

Figure 3.4 Structure of an MPLS Label Header

Label Value	EXP	S	TTL

An MPLS label header is composed of the following parts:

Label Value (20 bits)

The label value is an integer in the range 16 – 1048575. (Labels 0 – 15 are reserved by the IETF for special usage.) For signalled LSPs, the Foundry device dynamically assigns labels in the range 1024 – 499999. For static LSPs, you assign label values; Foundry recommends that you assign labels in the range 16 – 1023. You can still assign a label in the range 1024 – 499999 to a static LSP, but it may conflict with a label assigned to an existing signalled LSP. When there is such a conflict, the Foundry device displays an error message.

EXP field (3 bits)

The EXP field is designated for experimental usage. The Foundry device uses the EXP field to define a Class of Service (CoS) value for prioritizing packets travelling through an LSP. By default, the Foundry device copies the first two bits in an IP packet's ToS field to the to the EXP field. Optionally, you can manually specify a CoS value to be applied to the EXP field of all packets travelling through an LSP; this way all packets travelling through an LSP can be treated with the same priority as they travel the MPLS domain.

S (Bottom of Stack) field (1 bit)

An MPLS packet can be assigned multiple labels. If an MPLS packet has multiple labels, they are logically organized in a last-in, first-out ***label stack***. An LSR performs a pop or swap operation on the topmost label; that is, the most recently applied label in the stack. The Bottom of Stack field indicates whether this label is the last (oldest) label in the stack. If the label is the last one in the stack, the Bottom of Stack field is set to 1. If not, the Bottom of Stack field is set to 0.

A Foundry device acting as an LSR can perform one push, swap, or pop operation on an incoming MPLS packet. The Foundry device can accept MPLS packets that contain multiple labels, but only the topmost label is acted upon.

TTL field (8 bits)

The TTL field indicates the Time To Live (TTL) value for the MPLS packet. At the ingress LER, an IP packet's TTL value is copied to its MPLS TTL field. At each transit LSR hop, the MPLS TTL value is decremented by 1. If the MPLS TTL value reaches 0, the packet is discarded. At the egress LER, the incoming packet's MPLS TTL value is copied to the packet's IP TTL field, the IP TTL field is decremented by 1, and the checksum is recalculated. Optionally, you can configure the LSRs in an LSP not to decrement the MPLS TTL value at each hop, which causes the LSP to appear as a single IP hop.

Using MPLS in Traffic Engineering

Traffic engineering is the task of routing network traffic to avoid points of congestion and make efficient use of high bandwidth interfaces. When used as an application of MPLS, traffic engineering involves creating LSPs that make the best use of available network resources; that is, ***traffic-engineered LSPs***. This section explains the process involved in creating traffic-engineered LSPs.

Creating traffic-engineered LSPs involves the following tasks:

- Gathering information about the network
- Using the gathered information to select optimal paths through the network
- Setting up and maintaining the paths

For traffic-engineered signalled LSPs, Foundry devices can perform these tasks dynamically. For traffic-engineered static LSPs, you perform these tasks manually. Figure 3.5 illustrates the process that takes place to configure, establish, and activate traffic-engineered signalled and static LSPs.

Figure 3.5 How traffic-engineered LSPs are configured, established, and activated

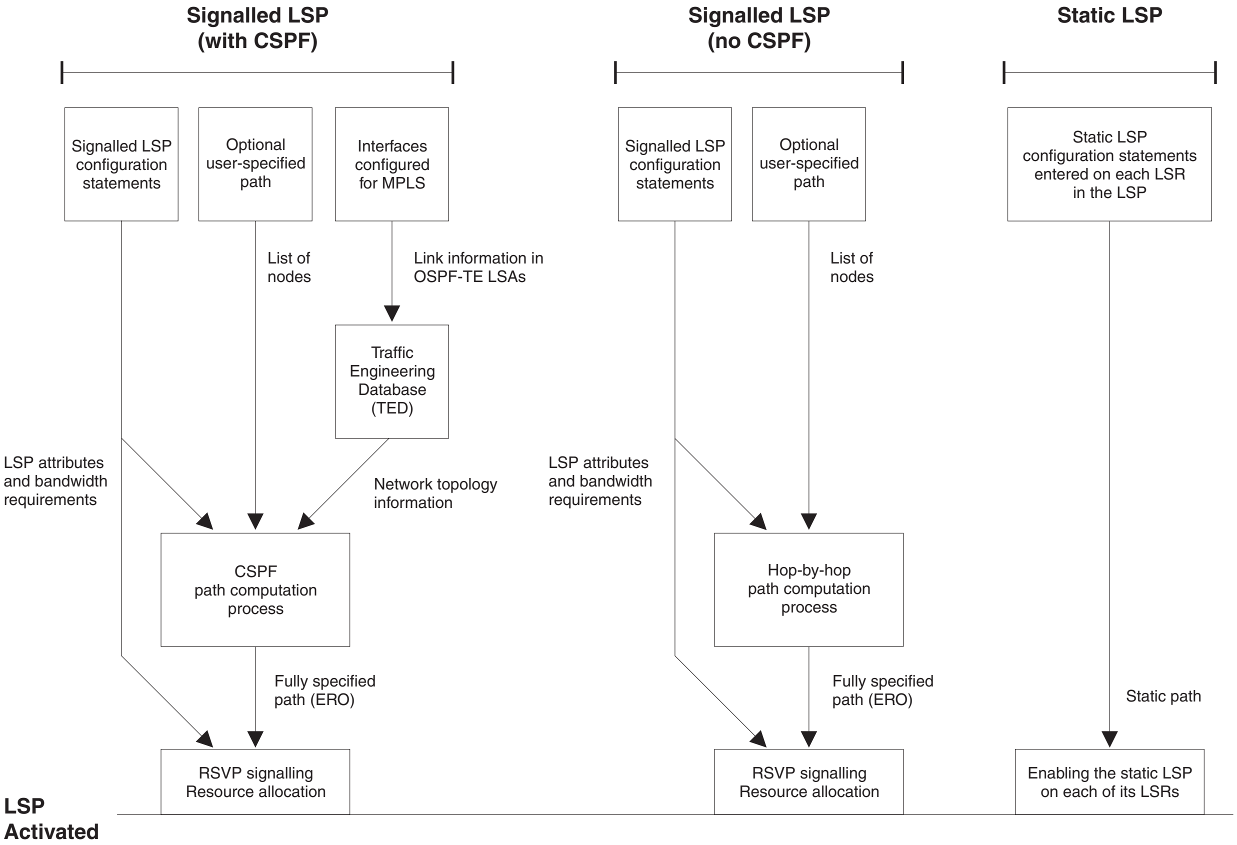
Signalled LSP (with CSPF)
Signalled LSP (no CSPF)
Static LSP
Signalled LSP configuration statements
Optional user-specified path
Interfaces configured for MPLS
Signalled LSP configuration statements
Optional user-specified path
Static LSP configuration statements entered on each LSR in the LSP
List of nodes
Link information in OSPF-TE LSAs
List of nodes
Traffic Engineering Database (TED)
LSP attributes and bandwidth requirements
Network topology information
LSP attributes and bandwidth requirements
CSPF path computation process
Hop-by-hop path computation process
Fully specified path (ERO)
Fully specified path (ERO)
Static path
RSVP signalling Resource allocation
RSVP signalling Resource allocation
Enabling the static LSP on each of its LSRs
LSP Activated

Traffic-engineered signalled LSPs are configured, established, and activated using the following process:

MPLS-enabled devices running OSPF can be configured to send out LSAs that have special extensions for traffic engineering. These LSAs, called ***OSPF-TE LSAs***, contain information about interfaces configured for MPLS. The OSPF-TE LSAs are flooded throughout the OSPF area. LSRs that receive the OSPF-TE LSAs place the traffic engineering information into a ***Traffic Engineering Database (TED)***, which maintains topology data about the nodes and links in the MPLS domain.

When you configure a signalled LSP, you specify the address of the egress LER, as well as optional attributes, such as the LSP's priority and bandwidth requirements. You can optionally specify a path of LSRs that the LSP must pass through on the way to the egress LER. When you enable the signalled LSP, the ***Constrained Shortest***

Path First (CSPF) process on the ingress LER uses this information to calculate a ***traffic-engineered path*** between the ingress and egress LERs.

CSPF is an advanced form of the Shortest Path First (SPF) process used by IGP routing protocols. The CSPF process on the ingress LER uses the configured attributes of the LSP, user-specified path (if there is one), and the information in the Traffic Engineering Database to calculate the traffic-engineered path, which consists of a sequential list of the physical interfaces that packets assigned to this LSP will pass through to travel from the ingress LER to the egress LER. The traffic-engineered path takes into account the network topology, available resources, and user-specified constraints. The traffic-engineered path calculated by CSPF may or may not be the same as the shortest path that would normally be calculated by standard IGP routing protocols.

CSPF is enabled by default for signalled LSPs, but can be disabled. When signalled LSPs are configured without CSPF, the shortest path from the ingress LER to the egress LER is calculated using standard hop-by-hop routing methods. If the LSP also is configured to use a user-specified path, the device calculates the shortest path between each LSR in the path. As with CSPF, the output of this process is a fully specified path of physical interfaces on LSRs.

The advantage of configuring signalled LSPs without CSPF is that it does not require LSRs to send out OSPF-TE LSAs. Since OSPF-TE LSAs have area flooding scope, the information in an LSR’s Traffic Engineering Database is relevant only to the OSPF area. Consequently, signalled LSPs that use CSPF can span only an OSPF area. Signalled LSPs that don’t use CSPF, because they do not rely on information in the TED, do not have this restriction; they can extend beyond an OSPF area.

Once the path for the LSP has been calculated, RSVP signalling then causes resources to be reserved and labels to be allocated on each LSR specified in the path. This may cause already existing, lower priority LSPs to be preempted. Once resources are reserved on all the LSRs in the path, the signalled LSP is considered to be ***activated***; that is, packets can be forwarded over it.

Static LSPs, in contrast, are configured manually on each LSR. No signalling protocol is used. On each LSR in a static LSP, you specify the labels to be applied by each LSR, as well as the inbound and outbound interfaces for labelled packets. When you enable the static LSP on each LSR in the path, it is activated.

To set up a traffic-engineered static LSP, you can analyze your network topology, looking for points of congestion or under-utilized resources. Using this information, you can calculate paths that make better use of network resources. Unlike signalled LSPs, static LSPs are never subject to preemption. Once enabled, static LSPs continue to be active until disabled. As with signalled LSPs that do not use CSPF, static LSPs can extend beyond an OSPF area.

The following sections provide additional information about the individual components of the process for activating traffic-engineered signalled LSPs, illustrated in Figure 3.5 on page 3-8.

OSPF-TE Link State Advertisements for MPLS Interfaces

Traffic engineering information is carried in OSPF traffic engineering (OSPF-TE) link state advertisements. OSPF-TE LSAs are Type 10 Opaque LSAs, as defined in RFC 2370. Type 10 Opaque LSAs have area flooding scope.

OSPF-TE LSAs have special extensions that contain information related to traffic engineering; these extensions are described in the IETF Internet Draft draft-katz-yeung-ospf-traffic-04.txt. The extensions consist of Type/Length/Value triplets (TLVs) containing the following information:

- Type of link (either point-to-point or multiaccess network)
- ID of the link (for point-to-point links, this is the Router ID of the LSR at the other end of the link; for multiaccess links, this is the address of the network’s designated router)
- IP address of the local interface for the link
- IP address of the remote interface for the link
- Traffic engineering metric for the link (by default, this is equal to the OSPF link cost)
- Maximum bandwidth on the interface
- Maximum reservable bandwidth on the interface

- Unreserved bandwidth on the interface
- Administrative group(s) to which the interface belongs

When configured to do so, the Foundry device sends out OSPF-TE LSAs for each of its MPLS-enabled interfaces. You can optionally specify the maximum amount of bandwidth that can be reserved on an interface, as well as assign interfaces to administrative groups. See "Setting Traffic Engineering Parameters for MPLS Interfaces" on page 3-25 for more information.

The following events trigger the Foundry device to send out OSPF-TE LSAs:

- Change in the interface's administrative group membership
- Change in the interface's maximum available bandwidth or maximum reservable bandwidth
- Significant change in unreserved bandwidth per priority level:
  - If for any priority level, the difference between the previously advertised unreserved bandwidth and the current unreserved bandwidth exceeds 5 percent of the maximum reservable bandwidth
  - Any changes while the total reserved bandwidth exceeds 95 percent of the maximum reservable bandwidth

In addition, OSPF-TE LSAs can be triggered by OSPF; for example, when an interface's link state is changed. When an interface is no longer enabled for MPLS, the device stops sending out OSPF-TE LSAs for the interface.

Traffic Engineering Database

An LSR's Traffic Engineering Database (TED) stores topology information about the MPLS domain. This topology information comes from the OSPF-TE LSAs that are flooded throughout the OSPF area. When an LSR receives OSPF-TE LSAs from neighboring LSRs, it places the traffic engineering information into its TED. In this way, each LSR in the OSPF area builds an identical topology database that reflects the traffic engineering constraints, bandwidth reservations, and administrative group memberships of the area's MPLS-enabled interfaces and the links connecting them.

The topology information in the TED is used by the CSPF process when it calculates traffic-engineered paths for signalled LSPs, as described in "How CSPF Calculates a Traffic-Engineered Path", below. You can display the contents of an LSR's TED; see "Displaying the Contents of the Traffic Engineering Database" on page 3-44.

LSP Attributes and Requirements Used for Traffic Engineering

In addition to the topology information in the TED, the Foundry device considers attributes and requirements specified in configuration statements for the LSP. The following user-specified parameters are considered when the device calculates a traffic-engineered path for a signalled LSP:

- Destination address of the egress LER
- Explicit path to be used by the LSP
- Class of Service (CoS) value assigned to the LSP
- Bandwidth required by the LSP
- Setup priority for the LSP
- Metric for the LSP
- Whether the LSP includes or excludes links belonging to specified administrative groups

See "Configuring Signalled LSP Parameters" on page 3-29 for more information on how to set these parameters.

How CSPF Calculates a Traffic-Engineered Path

Using information in the TED, as well as the attributes and requirements of the LSP, CSPF calculates a traffic-engineered path for the LSP by doing the following:

1. If more than one LSP needs to be enabled, select the LSP to calculate based on its setup priority and bandwidth requirement.

When multiple LSPs are enabled at once, such as when the Foundry device is booted, their paths are calculated one at a time, starting with the LSP that has the highest configured setup priority. If more than one LSP has the same setup priority, the LSP with the highest configured bandwidth requirement is calculated first.

2. Eliminate unsuitable links from consideration.

   The Foundry device examines the topology information stored in its TED and uses this information to eliminate links from consideration for the traffic-engineered path. Links are eliminated using the following procedure:

   - Eliminate any links that are not full duplex
   - Eliminate any links that do not have enough reservable bandwidth to fulfill the LSP's configured requirements
   - If the LSP has an **include** statement, eliminate any links that don't belong to administrative groups specified in the statement
   - If the LSP has an **exclude** statement, eliminate any links that belong to the administrative groups specified in the statement, as well as any links that don't belong to any administrative group

3. Using the remaining links, calculate the shortest path through the MPLS domain.

   Using the links that were not eliminated in the previous step, the device calculates the shortest path between the ingress and egress LERs. If the LSP is configured to use an explicit path, the device individually calculates the shortest path between each node in the path. See "Setting up Paths" on page 3-28 for more information on explicit paths.

   By default, the path calculated by CSPF can consist of no more than 255 hops, including the ingress and egress LERs. You can optionally change this maximum to a lower number. See "Limiting the Number of Hops the LSP Can Traverse" on page 3-33.

4. If multiple paths have the same cost, select one of them.

   The shortest path calculation performed in the previous step may result in multiple, equal-cost paths to the egress LER. In this case, the Foundry device chooses the path whose final node is the physical address of the destination interface.

   If more than one path fits this description, by default, the Foundry device chooses the path with the fewest hops. If multiple paths have this number of hops, the device chooses one of these paths at random. You can optionally configure the device to choose the path that has either the highest available bandwidth or the lowest available bandwidth. See "Specifying a Tie-Breaker for Selecting CSPF Equal-Cost Paths" on page 3-33.

The output of the CSPF process is a traffic-engineered path, a sequential list of the physical interfaces that packets assigned to this LSP pass through to reach the egress LER. Once the traffic-engineered path has been determined, RSVP signalling attempts to establish the LSP on each LSR in the path. See the next section, "How RSVP Establishes a Signalled LSP", for a description of how this works.

How RSVP Establishes a Signalled LSP

The traffic-engineered path calculated by CSPF consists of a sequential list of physical interface addresses, corresponding to a path from the ingress LER to the egress LER. Using this traffic-engineered path, RSVP establishes the forwarding state and resource reservations on each LSR in the path.

As with OSPF, special extensions for traffic engineering have been defined for RSVP. These extensions include the EXPLICIT_ROUTE, LABEL_REQUEST, LABEL, and RECORD_ROUTE objects as well as the Fixed Filter (FF) reservation style. These extensions are described in the IETF Internet Draft draft-ietf-mpls-rsvp-lsp-tunnel-08.txt.

The following diagram illustrates how RSVP establishes a signalled LSP:

Figure 3.6 How RSVP establishes a signalled LSP

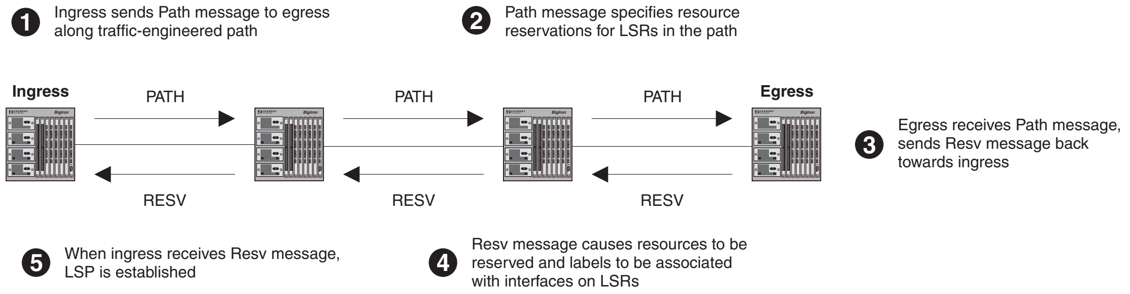

1 Ingress sends Path message to egress along traffic-engineered path

2 Path message specifies resource reservations for LSRs in the path

3 Egress receives Path message, sends Resv message back towards ingress

4 Resv message causes resources to be reserved and labels to be associated with interfaces on LSRs

5 When ingress receives Resv message, LSP is established

RSVP signalling for LSPs works as follows:

1. The ingress LER sends an RSVP Path message towards the egress LER.

   The Path message contains the traffic engineered path calculated by the CSPF process, specified as an EXPLICIT_ROUTE object (ERO). The Path message travels to the egress LER along the route specified in the ERO.

   The Path message also describes the traffic for which resources are being requested and specifies the bandwidth that needs to be reserved to accommodate this traffic. In addition, the Path message includes a LABEL_REQUEST object, which requests that labels be allocated on LSRs and tells the egress LER to place a LABEL object in the Resv message that it sends back to the ingress LER.

   Before sending the Path message, the ingress LSR performs ***admission control*** on the outbound interface, ensuring that enough bandwidth can be reserved on the interface to meet the LSP's requirements. Admission control examines the LSP's configured ***setup priority*** and ***mean-rate*** settings. For the LSP to pass admission control, the outbound interface must have reservable bandwidth at the LSP's setup priority level that is greater than the amount of bandwidth specified by the LSP's mean-rate setting. See "Admission Control, Bandwidth Allocation, and LSP Preemption", for more information and examples of this process.

2. The Path message requests resource reservations on the LSRs along the path specified in the ERO.

   If the LSP passes admission control, the ingress LER sends a Path message to the address at the top of the ERO list. This is the address of a physical interface on the next LSR in the path. As the ingress LER did, this LSR performs admission control to make sure the outbound interface has enough reservable bandwidth to accommodate the LSP.

   If the LSP passes admission control, the LSR then removes its address from the top of the ERO list and sends the Path message to the address now at the top of the ERO list. This process repeats until the Path message reaches the last node in the ERO list, which is the egress LER.

3. The egress LER receives the Path message and sends a Resv message towards the ingress LER.

   Resv messages flow upstream from the receiver of the Path message to the sender (that is, from the egress LER to the ingress LER), taking the exact reverse of the path specified in the ERO. In response to the LABEL_REQUEST object in the Path message, the Resv message from the egress LER includes a LABEL object. The LABEL object is used to associate labels with interfaces on the LSRs that make up the LSP.

4. As the Resv messages travel upstream, resources are reserved on each LSR.

   When an LSR receives a Resv message, it again performs admission control on the interface where the Resv message was received (that is, the interface that will be the outbound interface for packets travelling through the LSP). If the LSP still passes admission control, bandwidth is allocated to the LSP. The LSR allocates the amount of bandwidth specified by the LSP's mean-rate setting, using bandwidth available to its ***hold priority*** level. This may cause lower priority LSPs active on the device to be preempted.

   Once bandwidth has been allocated to the LSP, the LABEL object in the Resv message is used to associate labels with interfaces in the LSR's MPLS forwarding table. Figure 3.7 shows an example of how this works.

Figure 3.7 How the RSVP LABEL object associates a label with an interface in the MPLS forwarding table

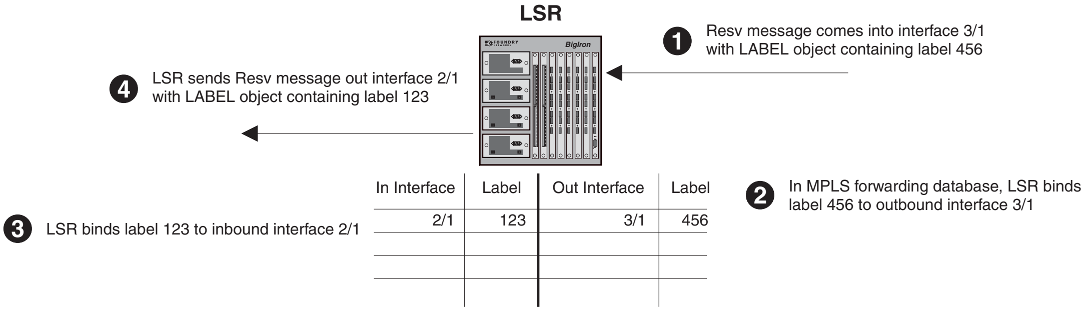

1 Resv message comes into interface 3/1 with LABEL object containing label 456

2 In MPLS forwarding database, LSR binds label 456 to outbound interface 3/1

3 LSR binds label 123 to inbound interface 2/1

4 LSR sends Resv message out interface 2/1 with LABEL object containing label 123

In Interface	Label	Out Interface	Label
2/1	123	3/1	456

In the example above, the LSR receives a Resv message on interface 3/1 from the downstream LSR in the ERO. The Resv message has a LABEL object containing label 456. After performing admission control and bandwidth allocation, the LSR adds an entry to its MPLS forwarding table for this LSP, associating label 456 with outbound interface 3/1.

The LSR then takes a label from its range of available labels (for example, 123) and places it in the LABEL object in the Resv message that it sends to the upstream LSR. In this example, the LSR sends the Resv message out interface 2/1 to the upstream LSR in the ERO. In its MPLS forwarding table for this LSP, the LSR associates label 123 with inbound interface 2/1.

This process repeats at each LSR until the Resv message reaches the ingress LER.

NOTE: To enable penultimate hop popping for the LSP, the LABEL object sent by the egress LER to the penultimate LSR contains a value of 3 (Implicit Null Label). This is an IETF-reserved label value that indicates to the penultimate LSR that it must pop the label of MPLS-encoded packets that belong to this LSP.

5. Once the Resv message reaches the ingress LER, and the process described in Step 4 takes place, the LSP is activated. At this point each LSR in the LSP has reserved resources, allocated labels, and associated labels with interfaces. The LSP is activated, and the ingress LER can assign packets to the LSP.

Refresh Messages

Once a signalled LSP is enabled at the ingress LER, the router persistently attempts to establish the LSP through periodic retries until the LSP is successfully established. To maintain the forwarding states and resource reservations on the routers in an LSP, Path and Resv messages are exchanged between neighboring LSRs at regular intervals. If these refresh messages are not received on the routers in the LSP, the RSVP forwarding states and resource reservations are removed. You can control how often the Path and Resv messages are sent, as well as how long the Foundry device waits before removing forwarding states and resource reservations. See “Setting RSVP parameters” on page 3-27 for more information.

Admission Control, Bandwidth Allocation, and LSP Preemption

When a Resv message is received on an LSR, admission control determines whether the LSP can be established, based on its configured priority. If an LSP passes admission control, bandwidth is allocated to the new LSP, possibly preempting already-existing LSPs that have lower priority.

An LSP’s priority consists of a ***setup*** priority and a ***hold*** priority. The setup priority is the priority for taking resources; the hold priority is the priority for holding resources. An LSP’s setup priority is considered during admission control, and its hold priority is considered when bandwidth is allocated to the LSP. The setup and hold priorities are expressed as numbers between 0 (highest priority level) and 7 (lowest priority level). An LSP’s setup priority must be lower than or equal to its hold priority. You can configure either of these values for an LSP; by default, an LSP’s setup priority is 7 and its hold priority is 0.

On an MPLS-enabled interface, a certain amount of bandwidth is allocated for usage by LSPs; this amount can be either the maximum available bandwidth on the interface (the default), or a user-specified portion. The amount of bandwidth an individual LSP can reserve from this pool of allocated bandwidth depends on two user-configured attributes of the LSP: the LSP’s priority, and the average rate of packets that can go through the LSP (the LSP’s ***mean-rate***).

- For an LSP to pass admission control, the bandwidth available to its setup priority level must be greater than the value specified by its mean-rate
- If an LSP passes admission control, the bandwidth specified by its mean-rate is allocated to the LSP, using bandwidth available to its hold priority level
- To allocate bandwidth to the new LSP, already-existing, lower-priority LSPs may be preempted

When setting up an LSP, the Foundry device actually performs admission control twice: when the Path message is received and when the Resv message is received. If the LSP passes admission control after the Resv message is received, bandwidth allocation and LSP preemption take place.

The following sections provide examples of how admission control, bandwidth allocation, and LSP preemption work.

Admission Control

Admission control examines the LSPs setup priority and mean-rate settings to determine whether the LSP can be activated. To pass admission control, the reservable bandwidth available at the LSP's setup priority level must be greater than the value specified by its mean-rate.

For example, if the maximum reservable bandwidth on an interface is 10,000 Kbits/second, and there are no currently active LSPs, the amount of reservable bandwidth on the interface for each priority level would be as follows:

Priority	Unreserved Bandwidth
0	10,000
1	10,000
2	10,000
3	10,000
4	10,000
5	10,000
6	10,000
7	10,000
Active LSPs: None	

The LSR receives a Resv message for an LSP that has a configured setup priority of 6 and a hold priority of 3. The mean-rate specified for this LSP is 1,000 Kbits/second. For priority level 6, up to 10,000 Kbits/second can be reserved. Since the configured mean-rate for this LSP is only 1,000 Kbits/second, the new LSP passes admission control.

Bandwidth Allocation

Once the LSP passes admission control, bandwidth is allocated to it. The bandwidth allocation procedure examines the LSP's hold priority and mean-rate settings. The amount of bandwidth specified by the mean-rate is allocated to the LSP, using reservable bandwidth available at the LSP's hold priority level.

In this example, the LSP's hold priority is 3 and mean-rate is 1,000 Kbits/second. On this interface, for priority level 3, up to 10,000 Kbits/second can be reserved. The amount of bandwidth specified by the mean-rate (1,000 Kbits/second) is allocated to the LSP.

After bandwidth is allocated to this LSP, the amount of unreserved bandwidth on the interface is reduced accordingly. In the example, the reservable bandwidth array for the interface now looks like this:

Priority	**Unreserved Bandwidth**
0	10,000
1	10,000
2	10,000
3	9,000
4	9,000
5	9,000
6	9,000
7	9,000
Active: LSP with setup 6, hold 3, mean-rate 1,000	

Given the bandwidth allocation above, if an LSP were established that had a setup priority of 3 and a mean-rate of 9,500 Kbits/second, it would not pass admission control, since only 9,000 Kbits/second is available at priority 3.

LSP Preemption

If there is not enough unallocated bandwidth on an interface to fulfill the requirements of a new LSP that has passed admission control, existing LSPs that have a lower priority may be preempted. When preemption occurs, bandwidth allocated to lower-priority LSPs is reallocated to the higher-priority LSP. LSP preemption depends on the bandwidth requirements and priority of the new LSP, compared to the bandwidth allocation and priority of already existing LSPs.

When LSP preemption is necessary, the Foundry device uses the following rules:

- Preempt existing LSPs that have lower priority than the new LSP
- If several existing LSPs have lower priority than the new LSP, preempt the LSP that has the lowest priority
- If two LSPs have equal priority, and one must be preempted, preempt the one with the higher bandwidth requirement
- Preempt the fewest number of LSPs necessary
- Never preempt static LSPs

In the example above, bandwidth has been allocated to an LSP that has a hold priority of 3 and a mean-rate of 1,000 Kbits/second. When a new LSP with a setup priority of 2, hold priority of 1, and mean-rate of 10,000 Kbits/second is established, admission control, bandwidth allocation, and LSP preemption work as follows:

1. Admission control: On the interface, there is 10,000 Kbits/second available to priority 2. The mean-rate for the new LSP is 10,000, so the LSP passes admission control; bandwidth can be allocated to it.
2. Bandwidth allocation: The hold priority for the new LSP is 1. On the interface, 10,000 Kbits/second is available to priority 1. This entire amount is allocated to the LSP.
3. LSP preemption: The first LSP had been using 1,000 Kbits/second of this amount, but its hold priority is only 3. Consequently, the first LSP is preempted, and its bandwidth allocation removed in order to make room for the new LSP.

Once this happens, the reservable bandwidth array for the interface looks like this:

Priority	Unreserved Bandwidth
0	10,000
1	0
2	0
3	0
4	0
5	0
6	0
7	0
Active: LSP with setup 2, hold 1, mean-rate 10,000 **Preempted:** LSP with setup 6, hold 3, mean-rate 1,000	

On this interface, the only LSP that could preempt the active LSP would be one that had a setup and hold priority of 0.

When several LSPs are candidates for preemption, the Foundry device normally preempts the one with the lowest priority. However, if preempting a higher priority LSP with a high bandwidth requirement would allow lower priority LSPs with lower bandwidth requirements to avoid preemption, then the higher priority LSP is preempted instead.

For example, consider an interface with 10,000 Kbits/second of reservable bandwidth, allocated to two active LSPs: one with a setup priority of 3, hold priority of 2, and mean-rate of 5,000 Kbits/second; and another with a setup priority of 4, hold priority of 3, and mean-rate of 2,500 Kbits/second. When an LSP with a setup priority of 1, hold priority of 0, and mean-rate of 7,500 Kbits/second is established, the following takes place:

1. Admission control: On the interface, there is 10,000 Kbits/second available to priority 1. The mean-rate for the new LSP is 7,500 Kbits/second, so the LSP passes admission control; bandwidth can be allocated to it.
2. Bandwidth allocation: The hold priority for the new LSP is 0. On the interface, 10,000 Kbits/second is available to priority 0. Of this amount, 7,500 Kbits/second is allocated to the new LSP.
3. LSP preemption: To reserve enough bandwidth for the new LSP, one of the active LSPs must be preempted. The LSP with hold priority 2 uses 5,000 Kbits/second, and the LSP with hold priority 3 uses 2,500 Kbits/second. Instead of preempting both LSPs, the device preempts the higher priority LSP and its allocation of 5,000 Kbits/second. This clears enough bandwidth to allow both the new LSP and the lower priority LSP to be active.

After preemption, the reservable bandwidth array for the interface looks like this:

Priority	Unreserved Bandwidth
0	2,500
1	2,500
2	2,500
3	0
4	0
5	0
6	0
7	0

Active: LSP with setup 1, hold 0, mean-rate 7,500

LSP with setup 4, hold 3, mean-rate 2,500

Preempted: LSP with setup 3, hold 2, mean-rate 5,000

NOTE: Static LSPs have a setup and hold priority of 0 and are never subject to preemption. When a static LSP is established, the bandwidth specified by its mean-rate is allocated to it, reducing the amount of bandwidth available to signalled LSPs. This may cause existing signalled LSPs to be preempted, using the rules described above.

Bandwidth allocations for static LSPs persist even when the amount of reservable bandwidth on an interface is reduced. If all the reservable bandwidth on an interface is allocated to static LSPs, no signalled LSPs can be activated.

Using Traffic Engineered LSPs Between IBGP Neighbors

Traffic-engineered LSPs can be used as shortcuts between IBGP neighbors in an Autonomous System. In this application, traffic headed for BGP destinations (destinations outside the AS) uses the traffic-engineered path specified by an LSP when travelling between IBGP neighbors to a BGP next hop address. The traffic-engineered path serves as an alternative to the IGP shortest path, which the BGP next hop traffic would normally take.

When a signalled or static LSP is activated, it becomes available to the BGP routing component as a potential route to its destination (the LSP’s egress LER). BGP always favors an LSP’s traffic-engineered path over the IGP shortest path. When traffic-engineered paths are available, BGP uses them automatically; no configuration is necessary.

Figure 3.8 shows an example of this kind of application.

Figure 3.8 Traffic engineered path between IBGP neighbors in an AS

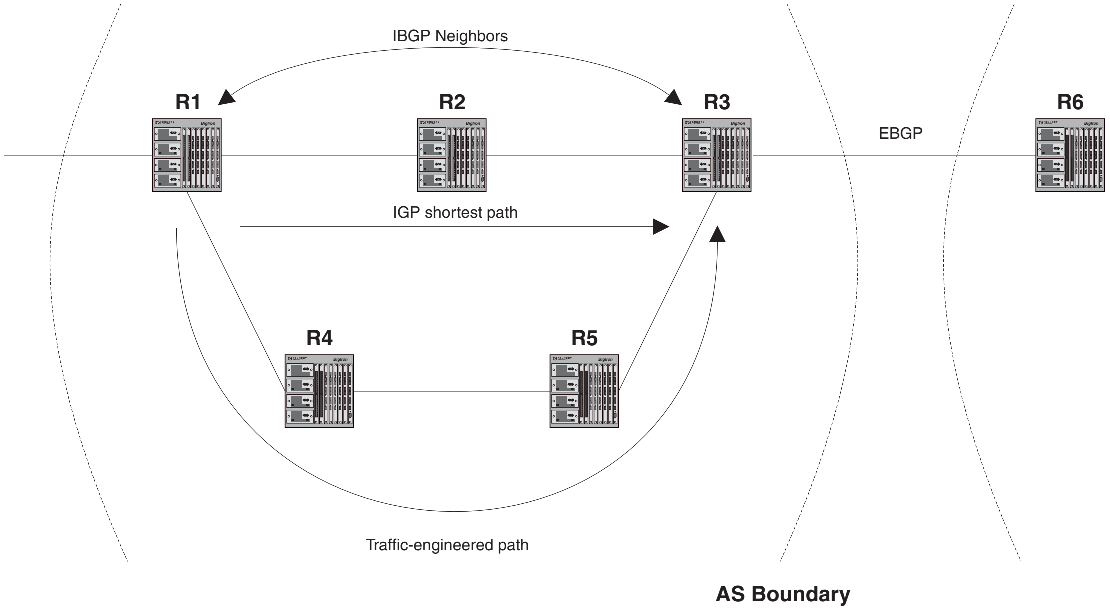

In the example above, BGP routers R1 and R3 are IBGP neighbors in an AS. On R1, LSP L1 specifies R3 as its destination. The LSP uses a traffic-engineered path through R4 and R5 to reach R3. This traffic-engineered path differs from the IGP shortest path, which goes through R2 to reach R3.

BGP installs LSP L1 as the outgoing "interface" for address prefixes for which R3 is the next hop in the main routing table. When traffic whose destination is R6 comes into R1, it is assigned to the LSP and forwarded along the traffic-engineered path to R3 (and then on to R6).

When LSP L1 activated, a 32-bit host route to R3 is placed into R1's ***MPLS routing table***. This host route consists of the following:

- Destination of the LSP – This is the address of the egress LER.
- MPLS tunnel interface port ID – This is a unique "virtual interface" identifier that corresponds to the LSP. This port ID allows the ingress LER to refer to the LSP in the same way it refers to physical, loopback, and VE ports. An LSP appears to be a physical interface directly connected to the destination node.
- User-configured metric for the LSP – This is a number in the range 1 – 65535, by default 1.

NOTE: You can display the contents of an LSR's MPLS routing table. See "Displaying the MPLS Routing Table" on page 3-51.

Only BGP uses information in the MPLS routing table. When BGP needs to resolve a BGP next hop address, it first performs a lookup in the MPLS routing table, comparing the BGP next-hop address to the LSP host routes in the table. If a matching LSP host route is found in the MPLS routing table, then BGP creates an ***MPLS tunnel route*** for every prefix associated with that BGP next hop address. BGP then injects these MPLS tunnel routes (corresponding to BGP destinations) into the main routing table (inet.0). In this way, the MPLS tunnel interface for the LSP is installed as the outgoing "interface" for address prefixes for which the egress LER is the BGP next hop.

In the example above, when BGP on R1 needs to resolve a next hop to R3, it looks in the MPLS routing table for a host route whose destination is R3. R3 is specified as the destination for the host route added to the table when LSP L1 was activated. Finding a match, BGP then creates MPLS tunnel routes for every prefix for which R3 is the BGP next hop, specifying LSP L1's port ID as the outgoing interface. BGP then injects these MPLS tunnel routes into main routing table (inet.0). The result is that all packets destined for a prefix for which R3 is the BGP next hop are forwarded along the path specified by LSP L1.

Whenever an LSP fails or is administratively disabled, at the ingress LER the LSP's host route is removed from the MPLS routing table, and an "LSP down" indication is given to BGP. BGP then removes from the main routing table all routes that use the disabled LSP. BGP may then attempt to reinstate the routes by using IGP routes instead of routes in the MPLS routing table.

See "LSP Between BGP Neighbors in an AS" on page 3-58 for a sample configuration.

Load Sharing for Equal-Cost MPLS Tunnel Routes

An LSP can be assigned a metric in the range 1 – 65535. A smaller metric value represents a lower cost; the default metric value for an LSP is 1. BGP includes this metric in the MPLS tunnel route it adds to the main routing table. When there are multiple, equal-cost MPLS tunnel routes to a destination, the traffic load is shared among them.

Figure 3.8 illustrates a configuration where two LSPs that have the same metric share the traffic load to a destination.

Figure 3.9 Load sharing for equal-cost MPLS tunnel routes to a destination

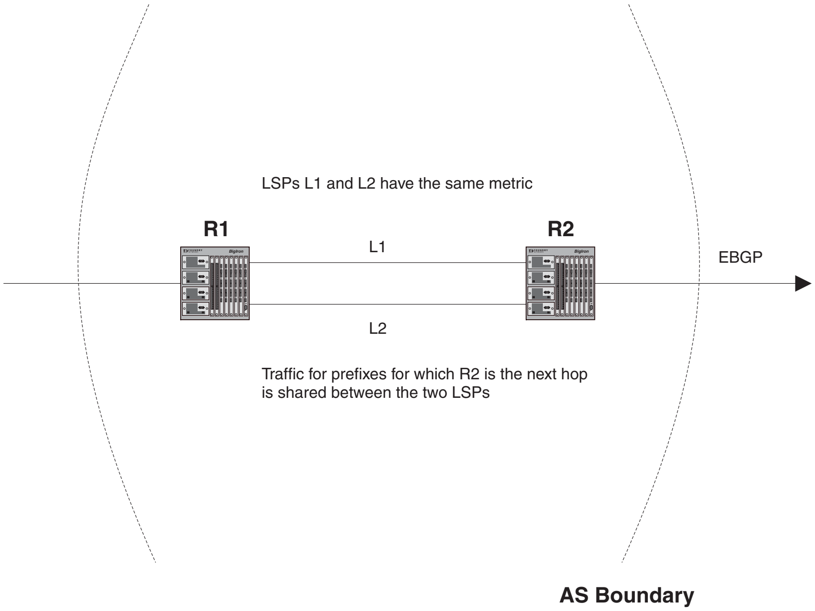
LSPs L1 and L2 have the same metric
R1
R2
L1
EBGP
L2
Traffic for prefixes for which R2 is the next hop
is shared between the two LSPs
AS Boundary

In this example, BGP router R1 has two LSPs, L1 and L2, to BGP router R2. Both LSPs have the same metric. In the main routing table, BGP installs both L1 and L2 as the outgoing "interfaces" for address prefixes for which R2 is the next hop. Packets for some of the prefixes are forwarded using the path specified by L1, while packets for the other prefixes are forwarded using the path specified by L2, thus splitting the load between the two paths.

The traffic load can be shared among LSPs with two different destinations if both destinations are equally favorable BGP next hops. Figure 3.8 shows an example of this kind of configuration.

Figure 3.10 Load sharing among LSPs with different tunnel end-points

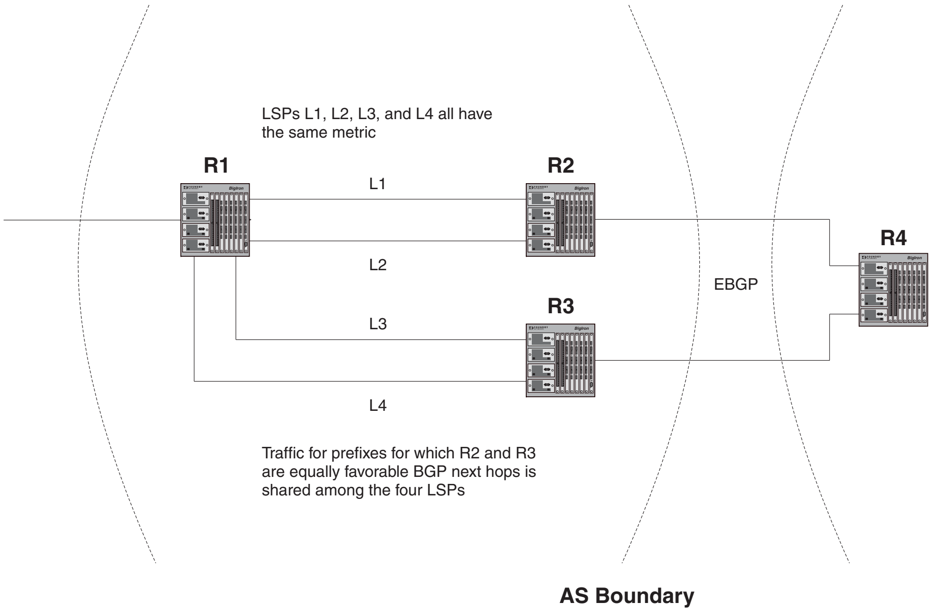
LSPs L1, L2, L3, and L4 all have the same metric
R1
R2
L1
L2
R4
EBGP
R3
L3
L4
Traffic for prefixes for which R2 and R3 are equally favorable BGP next hops is shared among the four LSPs
AS Boundary

In the example above, R1 has two LSPs, L1 and L2, to BGP router R2, as well as two LSPs, L3 and L4, to BGP router R3. All four LSPs have the same metric. If R2 and R3 are equally favorable next hops for some address prefixes, then the BGP component can install L1, L2, L3, and L4 as the outgoing “interfaces” for these address prefixes. In this way, the traffic load is shared not only among LSPs to the same tunnel end-point, but also among LSPs with different tunnel end-points.

The maximum number of equal-cost MPLS tunnel routes on the device is governed by the same rule for regular IP equal-cost paths; the default is four and the maximum configurable value is eight.

Using Aliases for Non-MPLS Routers

If a BGP next hop router is not MPLS capable, you can configure an LSP that treats the egress LER as an ***alias*** for the BGP next hop router. Traffic whose BGP next hop matches the alias is forwarded along the path specified by the LSP to the egress LER. Once this traffic reaches the egress LER, it can be forwarded to the BGP next hop router using the IGP shortest path. Figure 3.11 shows an example of this kind of configuration.

Figure 3.11 Using an alias for a non-MPLS BGP next hop router

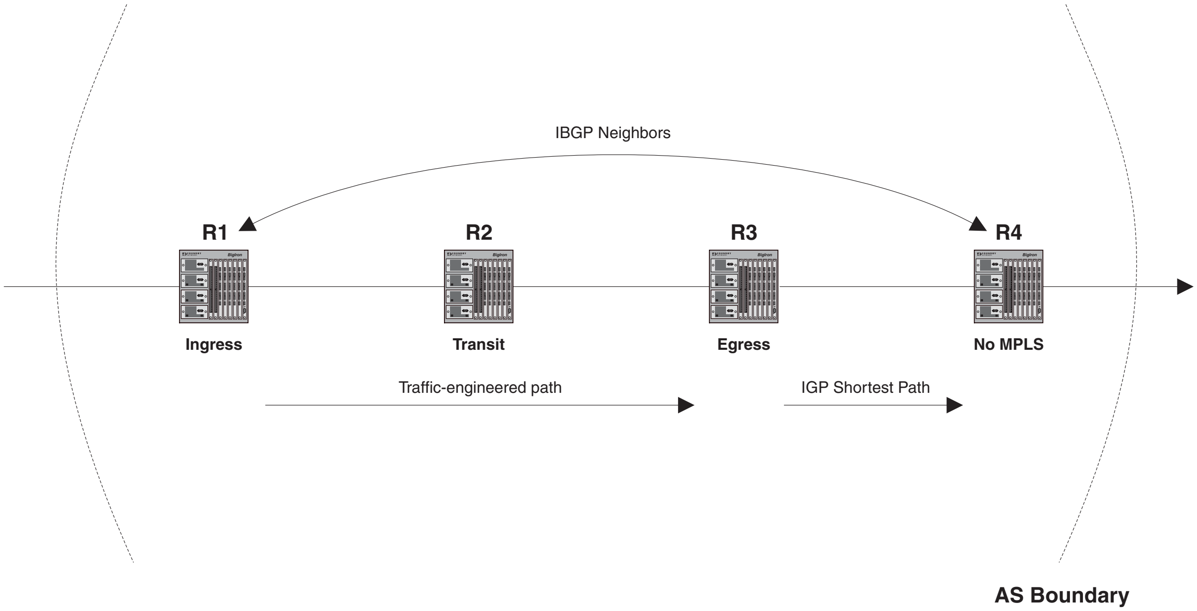

In this example, R1, R2, and R3 are MPLS capable, but R4 is not. An LSP is configured on R1 that specifies a path through R2 to reach R3. In the LSP's configuration, R4 is specified as an alias of R3. This means that when R1 receives packets bound for R4, it assigns them to the LSP, sending them to R3. When the packets exit the LSP at R3, they can be forwarded to R4 using standard hop-by-hop routing.

When the LSP is enabled, routes to both R3 and R4 are placed in R1's MPLS routing table, specifying that this LSP be used for packets going to either of these destinations. After BGP determines that R4 is the next hop address for a packet, it looks in the MPLS routing table for a match for this address. Finding a match for R4, BGP places the MPLS tunnel route into the main routing table.

In this example, BGP installs the LSP in R1's main routing table as the outgoing "interface" for address prefixes for which R4 is the next hop. Consequently, packets going from R1 to any prefixes for which R4 is the BGP next hop are always assigned to the LSP. When these packets reach R3 – the LSP's egress LER – they are forwarded to R4 using the IGP shortest path.

An alias can be an address prefix, thus allowing a single alias to represent multiple routers. Up to 12 aliases can be specified per LSP. See "LSP with an Alias for a Non-MPLS Destination" on page 3-61 for a sample configuration.

Using Traffic Engineered LSPs Within an AS

In addition to traffic destined outside an AS, Foundry devices can forward internal AS traffic into LSP tunnels; that is, MPLS traffic can be forwarded to IGP destinations.

In this application, you can configure a signalled LSP to serve as a shortcut between nodes in an AS. When an LSP is configured to be a shortcut, OSPF includes the LSP in its SPF calculation. If OSPF determines that the LSP shortcut is the best path to a destination, it installs a route into the IP routing table, specifying the LSP tunnel interface as the outbound interface, as well as the LSP's cost. Only LSPs configured to router IDs can be considered as shortcuts. If the LSP goes down or is administratively disabled, the LSP tunnel route is removed from the main routing table.

The cost of the LSP is the LSP's user-configured metric. If there is no user-configured metric for the LSP, the underlying IP cost of the LSP is used. For example, if the IP cost of the best underlying path between two routers is 2, and there is an LSP configured between these two routers, the LSP's cost would be 2. Once an LSP is used as a next hop for a destination, the cost of the LSP can be used for calculating other destinations that can use the LSP egress node as next hop. In this way, traffic for addresses downstream of the LSP egress node (including prefixes of the egress node) can use the LSP shortcut.

If OSPF is already using an LSP tunnel route to an Area Border Router (ABR), all inter-area routes via that ABR use the LSP as the next hop, provided there are no other better paths to the destination (that is, paths via other ABRs). An LSP to a destination outside an area is not used by OSPF in its calculation of inter-area routes.

Only signalled LSPs can be used as IGP shortcuts. A static LSP cannot be configured as an IGP shortcut. RSVP packets, used for establishing and maintaining signalled LSPs, are never forwarded into LSP tunnels.

See “Configuring the LSP to Serve as an IGP Shortcut” on page 3-34 for more information.

Configuring MPLS

This section explains how to set up MPLS on Foundry devices. It contains the following topics:

- "Enabling MPLS" on page 3-23
- "Setting Up Signalled LSPs" on page 3-27
- "Setting up Static LSPs" on page 3-35

Enabling MPLS

MPLS is disabled by default. To enable MPLS on a Foundry device, you perform the following steps:

1. Enable MPLS on the device
2. Enable MPLS on individual interfaces
3. Set global MPLS policy parameters (optional)
4. Set traffic engineering parameters for MPLS-enabled interfaces (optional)
5. Set RSVP parameters (optional)

Enabling MPLS on the Foundry Device

To enable MPLS on the device, enter the following commands:

```
NetIron> enable
NetIron# configure terminal
NetIron(config)# router mpls
```

Syntax: [no] router mpls

To disable MPLS on the device, use the **no** form of the command.

Enabling MPLS on Individual Interfaces

After you enable MPLS globally on the device, you can enable it either on all POS interfaces at once or on individual interfaces.

For example, to enable MPLS on all POS interfaces on the device:

```
NetIron(config-mpls)# mpls-interface all-pos
```

To enable MPLS only on interface 3/11:

```
NetIron(config-mpls)# mpls-interface pos 3/11
```

Syntax: [no] mpls-interface all-pos | <interface number>

The **all-pos** parameter enables MPLS on all POS interfaces.

Setting Global MPLS Policy Parameters

You can optionally set the following global MPLS policy parameters:

- Retry time
- Retry limit
- TTL propagation
- Administrative group names
- Whether the device sends out OSPF-TE LSAs for its MPLS-enabled interfaces

When you set these parameters, they are applied globally on all MPLS-enabled interfaces on the Foundry device.

Setting the Retry Time

When a signalled LSP is enabled, the ingress LER attempts to connect to the egress LER over the primary path specified in the LSP's configuration. If the connection is not successful, by default the ingress LER waits 30 seconds before attempting the connection again. You can configure the amount of time the ingress LER waits between connection attempts.

For example, to change the retry time to 45 seconds:

```
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# retry-time 45
```

Syntax: retry-time <seconds>

Setting the Retry Limit

If the ingress LER fails to connect to the egress LER in a signalled LSP, it will keep trying to make the connection indefinitely. You can set a limit for these connection attempts. After this limit is exceeded, the ingress LER stops trying to connect to the egress LER over the primary path.

If a secondary path is configured for the LSP, it is immediately activated once the primary path fails. After the secondary path is activated, the ingress LER continues to try to connect to the egress LER over the primary path, either up to the configured retry limit or indefinitely if no retry limit is set. If a connection over the primary path can be established, the secondary path is deactivated, and traffic for the LSP is again sent over the primary path.

To set the number of connection attempts to 20:

```
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# retry-limit 20
```

Syntax: retry-limit <number>

Once the connection is established, the retry counter is reset to zero. In the example above, if an LSP needs to be established again, the ingress LER will make 20 attempts to establish a connection to the egress LER.

Configuring TTL Propagation

In the MPLS label header, the TTL field indicates the Time To Live (TTL) value for an MPLS packet. At the ingress LER, an IP packet's TTL value is copied to its MPLS TTL field. At each transit LSR hop, the MPLS TTL value is decremented by 1. If the MPLS TTL value reaches 0, the packet is discarded.

At the MPLS router that pops the label (either the penultimate LSR or the egress LER), the incoming packet's MPLS TTL value is copied to the packet's IP TTL field, the IP TTL field is decremented by 1, and the checksum is recalculated. The result is that each LSR in the MPLS domain is counted as one hop. This is the default behavior.

Optionally, you can configure TTL propagation so that the entire MPLS domain appears as a single hop. In this case, the ingress LER places a value of 255 in the packet's MPLS TTL field. This value is decremented by 1 as the MPLS packet passes through each LSR in the MPLS domain. When the label is popped, the value in the MPLS TTL field is discarded, not copied to the packet's IP TTL field. The unlabeled IP packet's TTL is then decremented normally as it passes through the egress LER. This means that the packet's IP TTL is decremented only once from the time it enters the ingress LER to the time it exits the egress LER, making the MPLS domain appear as only one hop.

To configure TTL propagation so that the entire MPLS domain appears as a single hop, enter the following commands both on the ingress LER and the MPLS router that pops the label (either the penultimate LSR or the egress LER):

```
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# no propagate-ttl
```

Syntax: [no] propagate-ttl

When **no propagate-ttl** is configured, the ingress LER places a value of 255 into the packet's MPLS TTL field, regardless of the TTL value in the packet's IP header. When the MPLS label is popped, the LSR does not change the TTL value in the IP header.

NOTE: If you choose to configure TTL propagation in this way, it is important that you enter the **no propagate-ttl** command at ***both*** the ingress LER and the MPLS router that pops the label. If you omit the **no propagate-ttl** command at the MPLS router that pops the label, the value in the packet's MPLS TTL field would be copied into the packet's IP TTL field. This value could be as high as 255.

Establishing Administrative Group Names

Administrative groups, also known as resource classes or link colors, allow you to assign MPLS-enabled interfaces to various classes. When a Foundry device calculates the path for an LSP, it can take into account the administrative group to which a interface belongs; you can specify which administrative groups the Foundry device can include or exclude when making its calculation.

Up to 32 administrative groups can be configured on the Foundry device. You can refer to an administrative group either by name or number. To refer to an administrative group by name, you first establish a name for the group and associate the name with an administrative group number. This is done at the MPLS policy level.

For example, the following commands establish three administrative group names:

```
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# admin-group gold 30
NetIron(config-mpls-policy)# admin-group silver 20
NetIron(config-mpls-policy)# admin-group bronze 10
```

Syntax: [no] admin-group <name> <number>

The <number> can be from 0 – 31.

After you associate an administrative group name with a number, you can refer to it by name when assigning interfaces to the group or including or excluding the group from LSP calculations. See “Adding Interfaces to Administrative Groups” on page 3-26 and “Including or Excluding Administrative Groups from LSP Calculations” on page 3-33.

Enabling OSPF-TE LSAs for MPLS Interfaces

Information related to traffic engineering is carried in OSPF traffic engineering (OSPF-TE) LSAs. OSPF-TE LSAs have special extensions that contain information about an interface’s traffic engineering metric, bandwidth reservations, and administrative group memberships.

When an MPLS-enabled Foundry device receives an OSPF-TE LSA, it stores the traffic engineering information in its Traffic Engineering database (TED). The Foundry device uses information in the TED in its calculations when performing CSPF to determine a path for an LSP.

You can configure the Foundry device to send out OSPF-TE LSAs for all of its MPLS-enabled interfaces. To do this, enter the following commands:

```
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# traffic-engineering ospf
```

Syntax: [no] traffic-engineering ospf

By default, the Foundry device does not send out OSPF-TE LSAs for its MPLS-enabled interfaces. Since information in the TED is used to make path selections using CSPF, and information in the TED comes from OSPF-TE LSAs, you must enable the Foundry device to send out OSPF-TE LSAs if you want CSPF to perform constraint-based path selection.

See the next section, “Setting Traffic Engineering Parameters for MPLS Interfaces”, for information on the traffic engineering information carried in OSPF-TE LSAs.

Setting Traffic Engineering Parameters for MPLS Interfaces

When using constraints to determine a path for an LSP, the Foundry device takes into account information included in OSPF-TE LSAs. This information can be used to set up a path for a new LSP or to preempt an existing LSP so that an LSP with a higher priority can be established.

OSPF-TE LSAs include Type/Length/Value triplets (TLVs) containing the following information:

- Link type (either point-to-point or multiaccess network)
- Link ID (for point-to-point links, this is the Router ID of the LSR at the other end of the link; for multiaccess links, this is the address of the network’s designated router)
- IP address of the local interface
- IP address of the remote interface

- Traffic engineering metric for the link (by default, this is equal to the OSPF link cost)
- Maximum bandwidth on the interface
- Maximum reservable bandwidth on the interface
- Unreserved bandwidth on the interface
- Administrative group(s) to which the interface belongs

When configured to do so with the **traffic-engineering ospf** command, the Foundry device sends out OSPF-TE LSAs containing this information for each of its MPLS-enabled interfaces. Optionally, you can specify the maximum amount of bandwidth that can be reserved on an interface. In addition, you can assign interfaces to administrative groups.

Reserving Bandwidth on an Interface

An OSPF-TE LSA contains three TLVs related to bandwidth reservation:

- The Maximum Bandwidth TLV indicates the maximum outbound bandwidth that can be used on the interface. This TLV reflects the actual physical bandwidth of the interface (155M for OC-3, 622M for OC-12, or 2488M for OC-48). This TLV is not configurable by the user, although changing between an OC-3 module and an OC-12 module would cause an LSA to be issued that reflects the change in the interface's maximum available bandwidth.
- The Maximum Reservable Bandwidth TLV indicates the maximum bandwidth that can be reserved on the interface. By default, the Maximum Reservable Bandwidth is the same as the Maximum Bandwidth for the interface. You can optionally change the reservable bandwidth to an amount greater or less than the maximum available bandwidth of the interface.
- The Unreserved Bandwidth TLV indicates the amount of bandwidth not yet reserved on the interface. This TLV consists of eight octets, indicating the amount of unreserved bandwidth (in kbits per second) at each of eight priority levels. The octets correspond to the bandwidth that can be reserved with a hold priority of 0 through 7, arranged in increasing order, with priority 0 occurring at the start of the TLV, and priority 7 at the end of the TLV. The value in each of the octets is less than or equal to the maximum reservable bandwidth. The Unreserved Bandwidth TLV itself is not user-configurable, although it is affected by modifications to the reservable bandwidth on an interface, as well as changes to LSPs.

You can optionally change the amount of reservable bandwidth on an MPLS-enabled interface (that is, modify the value in the Maximum Reservable Bandwidth TLV in OSPF-TE LSAs sent out for the interface). To do this, enter commands such as the following:

```
NetIron(config-mpls)# mpls-interface pos 3/11
NetIron(config-mpls-interface)# reservable-bw 10000
```

Syntax: reservable-bw <number>

The reservable bandwidth is expressed in Kbits/sec. By default, the reservable bandwidth is the same as the maximum available bandwidth on the interface. If the amount of reservable bandwidth is greater than the maximum available bandwidth, then the link can be oversubscribed. If the reservable bandwidth is less than the maximum available bandwidth, then LSPs cannot reserve all physical bandwidth on the interface.

Changing the amount of reservable bandwidth on an interface causes the amount of unreserved bandwidth to be recalculated. In addition, it may cause an OSPF-TE LSA to be issued, as well as possibly pre-empt existing LSPs if bandwidth reservations can no longer accommodate them.

Adding Interfaces to Administrative Groups

You can place individual interfaces into administrative groups. Administrative groups, also known as resource classes or link colors, allow you to assign MPLS-enabled interfaces to various classes. For example, you can define a group called "gold" and assign high-bandwidth interfaces to it. When a Foundry device calculates the path for an LSP, it can take into account the administrative group to which a interface belongs. You can configure up to 32 administrative groups. By default, an interface does not belong to any administrative groups.

Administrative groups are numbered from 0 – 31. You can refer to an administrative group either by name or number. To refer to an administrative group by name, first create a name for the group and associate the name with an administrative group number. See "Establishing Administrative Group Names" on page 3-25 for information on how to do this.

To assign an MPLS-enabled interface to an administrative group called “gold”, enter commands such as the following:

```
NetIron(config-mpls)# mpls-interface pos 3/11
NetIron(config-mpls-interface)# admin-group gold
```

Syntax: admin-group <number> | <name> ...

The <number> can be from 0 – 31. The <name> must be a previously configured administrative group name.

An MPLS-enabled interface can belong to any number of administrative groups. For example, to assign an interface to group “gold” and group 31, enter commands such as the following:

```
NetIron(config-mpls)# mpls-interface pos 3/11
NetIron(config-mpls-interface)# admin-group gold 31
```

After you add interfaces to administrative groups, you can specify which groups can be included or excluded from LSP calculations. See “Including or Excluding Administrative Groups from LSP Calculations” on page 3-33.

Setting RSVP parameters

RSVP is automatically enabled when MPLS is enabled on the device. You can optionally configure the following RSVP parameters:

- refresh interval
- refresh multiple

Setting the Refresh Interval

To maintain path states and resource reservations on the routers in an LSP, RSVP Path and Resv messages are sent at regular intervals. Path messages flow downstream in an LSP, from the ingress LER towards the egress LER. Resv messages flow upstream, in the reverse direction of Path messages.

You can control how often the Path and Resv messages are sent by setting the refresh interval. By default, the refresh interval is 30 seconds. You can set the refresh interval to between 0 – 2147483 seconds.

To set the refresh interval to 20 seconds

```
NetIron(config-mpls)# rsvp
NetIron(config-mpls-rsvp)# refresh-interval 20
```

Syntax: refresh-interval <seconds>

Setting the Refresh Multiple

If refresh messages are not received, RSVP path states and resource reservations are removed from the routers in an LSP. By default, the Foundry device waits the length of 3 refresh intervals; if no refresh message is received by the end of that time, the path state or resource reservation is removed.

The refresh multiple is the number of refresh intervals that must elapse without a refresh message before a path state or resource reservation times out. By default, the refresh multiple is 3 intervals. You can set the refresh multiple to between 0 – 65535 intervals.

To set the refresh multiple to 5 intervals:

```
NetIron(config-mpls)# rsvp
NetIron(config-mpls-rsvp)# refresh-multiple 5
```

Syntax: refresh-multiple <intervals>

Setting Up Signalled LSPs

An LSP consists of an actual path of MPLS routers through a network, as well as the characteristics of the path, including bandwidth allocations and routing metrics. There are two kinds of LSPs: signalled and static. Signalled LSPs are configured at the ingress LER. When you enable a signalled LSP, RSVP causes resources to be allocated on the other routers in the LSP.

Configuring a signalled LSP consists of the following tasks:

- Specifying a path for the LSP to follow (optional)
- Setting parameters for the signalled LSP
- Specifying which packets are to be forwarded along the LSP (optional)

Setting up Paths

A ***path*** is a list of router hops that specifies a route across an MPLS domain. Once you create a path, you can create signalled LSPs that refer to the path. Paths are configured separately from LSPs so that a path may be specified once and then used by several LSPs that refer to the path by name. An LSP may specify a primary and one or more redundant paths.

A path is always configured at the ingress LER and assumes that the ingress LER is the beginning of the path. A path can contain any number of ***nodes***, which correspond to MPLS-enabled routers in the network. Each node has one attribute: whether it is ***strict*** or ***loose***. A strict node means that the router must be directly connected to the preceding node. A loose node means that there can be other routers in between.

Creating a path is not absolutely necessary when configuring an LSP. If you configure a signalled LSP without naming a path, CSPF uses only information in the Traffic Engineering Database (TED), as well as the user-configured attributes and requirements of the LSP to calculate the path. See "How CSPF Calculates a Traffic-Engineered Path" on page 3-10 for more information. If the LSP has been configured not to use CSPF, the path between the ingress and egress LERs is determined using standard hop-by-hop routing methods, as if the path consisted of a single loose node.

The following commands set up a path called sf_to_sj that has four nodes.

```
NetIron(config-mpls)# path sf_to_sj
NetIron(config-mpls-path)# strict 216.150.1.1
NetIron(config-mpls-path)# strict 216.150.1.2
NetIron(config-mpls-path)# loose 64.1.1.1
NetIron(config-mpls-path)# strict 64.100.1.1
NetIron(config-mpls-path)# exit
```

Syntax: [no] path <path name>

Syntax: [no] strict | loose <ip address>

The path is assumed to start from the local node. You specify the nodes in order from ingress to egress. Specifying the local node itself as the first node in the path is optional. Further, the final node does not necessarily have to be the egress LER in the LSP. (The egress LER is specified at the LSP configuration level with the **to** command.) If the final node in the path differs from the egress LER, the hop between the final node in the path and the egress LER is treated as a hop to a loose node; that is, standard IP routing is used to determine the path between the final node and the egress LER.

The IP address defines an LSR and can be any interface address or a loopback interface address on the LSR.

The **strict** and **loose** parameters are relative to the preceding node. In the sf_to_sj path defined above, LSR 216.150.1.2 is a strict node; it must be directly connected to LSR 216.150.1.1. LSR 64.1.1.1 is a loose node; this means there can be other routers between LSR 216.150.1.2 and 64.1.1.1. When specifying a strict node, you should make sure that the LSR is actually directly connected to the preceding node.

Modifying a Path

Once you have created a path, you can insert or delete nodes from it. For example, to delete a node from the sf_to_sj path defined above:

```
NetIron(config-mpls)# path sf_to_sj
NetIron(config-mpls-path)# delete loose 64.1.1.1
NetIron(config-mpls-path)# exit
```

Syntax: delete strict | loose <ip address>

To insert a node into a path:

```
NetIron(config-mpls)# path sf_to_sj
NetIron(config-mpls-path)# insert strict 216.150.1.1 before 216.150.1.2
NetIron(config-mpls-path)# exit
```

Syntax: insert strict | loose <ip address> before <ip address>

The **insert** command allows a new node to be inserted in front of an existing node within the path. In this example, the **insert strict 216.150.1.1 before 216.150.1.2** command assumes that 216.150.1.2 is already in the path and inserts 216.150.1.1 before it.

NOTE: When you modify a path, the changes are not carried over to active LSPs that refer to the path until the LSPs are deactivated and reactivated. For example, path sj_to_sf may be used by an LSP called lsp1. After lsp1 has been activated, any changes to path sj_to_sf do not cause the route followed by lsp1 to be modified. To get the LSP to use the modified path, you must deactivate and then reactivate lsp1.

Deleting a Path

To delete an entire path from the LSR's configuration, enter a command such as the following:

```
NetIron(config-mpls)# no path sf_to_sj
```

Syntax: no path <path name>

Configuring Signalled LSP Parameters

Once you have configured a path, you can configure signalled LSPs that refer to it. An LSP's configuration can specify not only the path that label-switched packets follow in a network, but also the characteristics of the path, the resources allocated along the path, and actions applied to the packets by the ingress or egress LERs.

You can perform the following tasks when configuring a signalled LSP:

- Creating the LSP
- Specifying an egress LER for the LSP
- Specifying a primary path for the LSP (optional)
- Configuring secondary or hot-standby paths for the LSP (optional)
- Setting aliases for the egress LER (optional)
- Setting a Class of Service (CoS) value for the LSP (optional)
- Allocating bandwidth to the LSP (optional)
- Configuring the setup and hold priority for the LSP (optional)
- Setting a metric for the LSP (optional)
- Including or excluding administrative groups from LSP calculations (optional)
- Limiting the number of hops the LSP can traverse (optional)
- Specifying a tie-breaker for selecting CSPF equal-cost paths (optional)
- Disabling the Record-Route function (optional)
- Disabling CSPF path calculations (optional)
- Configuring the LSP to serve as an IGP shortcut to the egress LER (optional)
- Enabling the LSP
- Disabling the LSP

Creating an LSP

To create a signalled LSP and enter the LSP configuration level, enter commands such as the following:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)#
```

Syntax: [no] lsp <name>

Specifying the Egress LER

Each LSP requires one and only one egress LER. The egress LER is the router from which packets exit the MPLS domain in this LSP. After the LSP is successfully established, the address of the egress LER is installed as an internal host route on the ingress LER, allowing the ingress LER to direct BGP next-hop traffic into the LSP. The destination address does not necessarily have to be the final node in the primary path specified for the LSP. If the final node in the path differs from the destination address, the hop between the final node in the path and the egress LER is treated as a loose hop.

To specify 64.100.1.1 as the address of the egress LER for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# to 64.100.1.1
```

Syntax: to <ip address>

The egress LER is the only required parameter in an LSP. All other parameters are optional.

Specifying the Primary Path for an LSP

The primary path is the route that packets travel when going through an LSP. You can specify a user-defined path or no path at all. See “Setting up Paths” on page 3-28 for information on defining a path. Once the LSP is enabled, the ingress LER attempts to signal the other LSRs in the path so that resources can be allocated to the LSP. If you do not specify a primary path, the path used in the LSP is the shortest path to the egress LER, as determined from standard IP routing methods, or CSPF if it is enabled.

To specify the sf_to_sj path as the primary path for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# primary-path sf_to_sj
```

Syntax: primary-path <path name>

Configuring Redundant Paths for an LSP

A signalled LSP has a primary path, which is either user-defined or computed by the ingress LER. Optionally, you can configure one or more redundant paths to serve as a backup. If the primary path fails, traffic for the LSP can be forwarded over the redundant path. When no redundant path is configured for the LSP, if the primary path fails, the ingress LER automatically attempts to compute a new path to the egress LER, establish the new path, and then redirect traffic from the failed path to the new path.

Configuring a redundant path allows you to exercise greater control over the rerouting process than if the ingress LER simply calculated a new path to the egress LER. When a redundant path is configured, if the primary path fails, the ingress LER attempts to establish the redundant path. As with the primary path, a redundant path follows an explicit route of loose or strict hops.

By default, the redundant path is established only when the primary path fails. You can optionally configure a redundant path to operate in ***hot-standby*** mode. A hot-standby path is established at the same time the primary path in the LSP is established. Resources are allocated to the hot-standby path, although no packets for the LSP are sent over the hot-standby path until the primary path fails. When the primary path fails, the already-established hot-standby path immediately takes over from the primary path. Since the hot-standby path is already active, service outages that can arise from the process of signaling and establishing a new path are eliminated.

After the redundant path has been activated, the ingress LER continues to try to connect to the egress LER over the primary path, either indefinitely or up to the configured retry limit. If a connection over the primary path can be established, the redundant path is deactivated, and traffic for the LSP is again sent over the primary path. Once the primary LSP becomes available again, the redundant path is torn down; if the path is a hot-standby path, it reverts to its backup status.

You can configure multiple redundant paths. When the primary path fails, the ingress LER attempts to establish a connection to the egress LER using the first redundant path configured for the LSP. If a connection cannot be established using the first redundant path, the second redundant path is tried, and so on. If a connection cannot be established after trying each redundant path in the configuration, the first redundant path is tried again, and the process repeats.

To configure a secondary path, first create a path, as described in "Setting up Paths" on page 3-28. After you create the path, you can specify that it is to be used as a redundant path. For example, the following commands cause a path called alt_sf_to_sj to be used if the primary path in LSP tunnel1 fails.

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# secondary-path alt_sf_to_sj
NetIron(config-mpls-lsp-sec-path)#
```

Syntax: secondary-path <path name>

Issuing the **secondary-path** command enters the secondary path configuration level. From this level, you can specify that this path is to operate in hot standby mode. For example:

```
NetIron(config-mpls-lsp-sec-path)# standby
```

Syntax: standby

Once the LSP is enabled, both the primary and hot-standby paths are activated, although packets are directed over only the primary path.

NOTE: At the secondary path level, you can configure separate values for the following parameters: Class of Service (CoS), setup and hold priority, bandwidth allocations, and inclusion or exclusion of interfaces in administrative groups. If you do not configure these parameters at the secondary path level, the secondary path will use the default values for these parameters.

Creating Aliases for the Egress LER

Traffic whose BGP next hop is the egress LER in the LSP is always forwarded along the LSP. In addition, you can configure one or more ***aliases*** for the egress LER, so that traffic whose BGP next hop matches one of the aliases is also forwarded along the LSP. Creating aliases is useful in situations where a BGP next-hop router is not MPLS capable. Traffic for the alias can be sent along the LSP to the egress LER, and then to the non-MPLS router using Layer 3 forwarding.

For example, to install 1.1.0.0/16 as an alias for the egress LER in LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# install 1.1.0.0/16
```

Syntax: install <ip address/mask>

Traffic whose BGP next hop matches 1.1.0.0/16 is then forwarded along the LSP to the egress LER. When it arrives at the egress LER, it can be forwarded to the actual BGP next hop. You can specify up to 12 aliases for the egress LER.

Setting a Class of Service Value for the LSP

The 3-bit EXP field in the MPLS header can be used to define a Class of Service (CoS) value for packets travelling through the LSP. The CoS value is used to specify a priority for MPLS packets.

There are two ways a CoS value can be applied to packets travelling through an LSP:

- Use the Type of Service (ToS) field in the IP header. This is the default behavior. Specifically, the Foundry device copies the first three bits in the packet's ToS field to the to the CoS (EXP) field in the MPLS header. The ToS value maps to one of the four priority queues on the Foundry device.
- Manually set a CoS value for the LSP. The CoS value that you set is applied to the CoS (EXP) field in the MPLS header of all packets entering this LSP. This way all packets travelling through an LSP can be treated with the same priority as they travel the MPLS domain. You can assign the LSP a CoS value from 0 to 7.

To assign a CoS value of 7 (highest priority) to all packets traveling through LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# cos 7
```

Syntax: cos <number>

The CoS value can be an integer from 0 – 7. When the label is popped, the CoS value in the MPLS header is discarded; it is not copied back it the IP ToS field. The MPLS CoS value is used only for determining priority within an MPLS domain.

Allocating Bandwidth to an LSP

You can specify the amount of bandwidth allocated to an LSP, including the maximum and average rate of packets that will be travelling through the LSP. Allocating bandwidth to an LSP allows the LSRs to determine how much bandwidth may be consumed by the LSP, as well as how much available bandwidth resources can be advertised via OSPF-TE LSAs.

You can specify that data be sent over the LSP at an average of <mean-rate> kbits per second. If necessary, data can be sent at <max-rate> kbits per second, as long as the burst sent at the maximum rate contains no more than <max-burst> bytes.

To set the maximum rate of packets that can go through an LSP (in Kbits/sec):

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# traffic-eng max-rate 20
```

Syntax: traffic-eng max-rate <rate>

To set the average rate of packets that can go through an LSP (in Kbits/sec):

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# traffic-eng mean-rate 10
```

Syntax: traffic-eng mean-rate <rate>

To set the maximum size (in bytes) of the largest burst the LSP can send at the maximum rate:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# traffic-eng max-burst 10
```

Syntax: traffic-eng max-burst <bytes>

Configuring a Signalled LSP’s Priority

You can specify a priority for each signalled LSP for which this is the ingress LER. The priority determines the relative importance of the LSP during setup or preemption. The priority for an LSP has two components: the setup priority and the hold priority.

When multiple LSPs are enabled at the same time, such as when the device is booted, LSPs that have a higher setup priority are enabled before LSPs that have a lower setup priority.

If an LSP is assigned a high setup priority, it may preempt an LSP that is already established, causing resources assigned to the lower priority LSP to be diverted to the higher priority LSP. The hold priority specifies how likely an established LSP is to give up its resources to another LSP. To be preempted, an LSP must have a lower hold priority than the preempting LSP’s setup priority. In addition, an established LSP can be preempted by a higher priority LSP only if it would allow the higher priority LSP to be established successfully.

To configure LSP tunnel1 with a setup priority of 6 and hold priority of 1.

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# priority 6 1
```

Syntax: priority <setup-priority> <hold-priority>

Possible values are 0 (highest priority) through 7 (lowest priority). An LSP’s setup priority must be lower than or equal to its hold priority. By default, an LSP’s setup priority is 7 and its hold priority is 0.

Assigning a Metric to the LSP

You can assign a metric to the LSP, which can be used by routing protocols to determine the relative preference among several LSPs towards a given destination.

To assign a metric of 5 to LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# metric 5
```

Syntax: metric <number>

The metric can be a number between 1 – 65535. By default, all LSPs have a metric of 1. A lower metric is preferred over a higher one. If there are multiple LSPs to the same destination LSR, and they share the same metric, the traffic load is shared among them.

Including or Excluding Administrative Groups from LSP Calculations

Administrative groups, also known as resource classes or link colors, allow you to assign MPLS-enabled interfaces to various classes. When a Foundry device calculates the path for an LSP using CSPF, it takes into account the administrative group to which a interface belongs; you can specify which administrative groups the Foundry device can include or exclude when making its calculation.

For example, to include interfaces in either administrative group “gold” or “silver” in the path calculations for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# include-any gold silver
```

Syntax: [no] include-any <groups>

The value specified for <groups> can be one or more valid administrative group names or numbers. In this example, the Foundry device includes any of the interfaces that are members of groups “gold” or “silver” when calculating the path for this LSP. Only those interfaces in the “gold” or “silver” groups are considered for the LSP. Interfaces that are not part of these groups, as well as interfaces that are not part of any group, are eliminated from consideration.

To exclude interfaces in either administrative group “gold” or “silver” when the path for LSP tunnel1 is calculated:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# exclude-any gold silver
```

Syntax: [no] exclude-any <groups>

In this example, the Foundry device excludes any of the interfaces that are members of groups “gold” or “silver” when calculating the path for this LSP. Only interfaces that are not part of either group can be considered for the LSP.

To specify that an interface must be a member of both the “gold” or “silver” administrative groups in order to be included in the path calculations for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# include-all gold silver
```

Syntax: [no] include-all <groups>

In this example, an interface must be a member of all the groups specified in the **include-all** command in order to be considered for the LSP. Any interface that is not a member of all the groups is eliminated from consideration.

Limiting the Number of Hops the LSP Can Traverse

By default, the path calculated by CSPF can consist of no more than 255 hops, including the ingress and egress LERs. You can optionally change this maximum to a lower number.

For example, to limit CSPF to choosing a path consisting of no more than 20 hops for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# hop-limit 20
```

Syntax: [no] hop-limit <number>

The number of hops can be from 0 – 255.

Specifying a Tie-Breaker for Selecting CSPF Equal-Cost Paths

CSPF may calculate multiple, equal-cost paths to the egress LER. When this happens, the Foundry device chooses the path whose final node is the physical address of the destination interface. If more than one path fits this description, by default, the Foundry device chooses the path with the fewest hops. If multiple paths have this number of hops, the device chooses one of these paths at random. You can optionally configure the device to choose the path that has either the highest available bandwidth or the lowest available bandwidth.

For example, the following commands cause CSPF to select the path with the highest available bandwidth when choosing among equal-cost paths calculated for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# tie-breaking least-fill
```

Syntax: [no] tie-breaking least-fill | most-fill | random

The **least-fill** parameter causes CSPF to choose the path with the highest available bandwidth (that is, the path with the least utilized links).

The **most-fill** parameter causes CSPF to choose the path with the lowest available bandwidth (that is, the path with the most utilized links).

The **random** parameter causes CSPF to choose the path randomly from among the equal-cost paths. This is the default.

Disabling the Record Route Function

The RSVP RECORD_ROUTE object (RRO) allows an LSP's path to be recorded. An RRO consists of a series of subobjects that can contain the addresses of the LSRs in the path. This information can be viewed with the **show mpls lsp detail** command. The path information is recorded in the RRO by default, but you can disable path recording.

To disable path recording in the RRO:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# no record
```

Syntax: [no] record

Disabling CSPF Path Calculations

By default, CSPF is enabled for signalled LSP calculations. That is, if the Foundry device receives OSPF-TE LSAs, it places the traffic engineering information from them in its Traffic Engineering Database (TED). When the Foundry device is the ingress LER for the LSP, it uses the information in the TED to help determine a path for the LSP. If all nodes in your network are not capable of sending out OSPF-TE LSAs, you may want to disable CSPF for the LSP.

To disable constraint-based path selection for LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# no cspf
```

Syntax: no cspf

Configuring the LSP to Serve as an IGP Shortcut

A signalled LSP can serve as an IGP shortcut between nodes in an AS. When an LSP is configured to be an IGP shortcut, OSPF includes the LSP in its SPF calculation. If OSPF determines that the LSP shortcut is the best path to a destination, it installs a route into the IP routing table, specifying the LSP tunnel interface as the outbound interface, as well as the LSP's cost. Only LSPs configured to router IDs can be considered as shortcuts. If the LSP goes down or is administratively disabled, or the **shortcuts ospf** command is removed from its configuration, the LSP tunnel route is removed from the main routing table.

If OSPF is already using an LSP tunnel route to an Area Border Router (ABR), all inter-area routes via that ABR use the LSP as the next hop, provided there are no other better paths to the destination (that is, paths via other ABRs). An LSP to a destination outside an area is not used by OSPF in its calculation of inter-area routes.

RSVP packets, used for establishing and maintaining signalled LSPs, are never forwarded into LSP tunnels.

If an LSP is configured to use a primary path and one or more secondary paths, and the LSP is being used as a shortcut to IGP destinations, the primary path is used for those destinations. If the primary path fails, the secondary path is used for those destinations.

To specify that an LSP be used as a shortcut for IGP destinations, enter commands such as the following:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# shortcuts ospf
```

Syntax: [no] shortcuts ospf

Enabling a Signalled LSP

After you set the parameters for the signalled LSP, you can enable it. Enabling the LSP causes the path to be set up and resources reserved on the LSRs in the LSP's primary path. Enabling the LSP is the final step in configuring it.

To enable LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# enable
```

Syntax: enable

Disabling an LSP

Disabling an LSP de-activates it, but does not remove the LSP from the device's configuration. (To remove the LSP from the device's configuration, use the **no lsp <name>** command.) To make changes to an active LSP, first disable the LSP, modify parameters on the LSP, and then enable the LSP.

To disable LSP tunnel1:

```
NetIron(config-mpls)# lsp tunnel1
NetIron(config-mpls-lsp)# disable
```

Syntax: disable

Assigning Packets to the LSP

When a packet enters the MPLS domain, the ingress LER classifies it and determines to which LSP (if any) the packet should be assigned. By default, packets whose BGP next-hop address is the egress LER (or one of its aliases) are always assigned to the LSP. In addition, you can cause packets whose destination address matches a specified IP address prefix to be assigned to an LSP.

For example, to assign packets matching prefix 30.1.1.0/24 to LSP tunnel1:

```
NetIron(config-mpls)# ip route 30.1.1.0/24 lsp tunnel1
```

Syntax: ip route <dest-ip-addr>/<mask-bits> lsp <name> | static-lsp <name>

Assigning packets to an LSP in this way is similar to installing a static route. Instead of specifying an Ethernet, POS, or VE interface as the next hop, you specify the name of a static LSP; packets matching the specified prefix are assigned to the LSP.

Setting up Static LSPs

A static LSP differs from a signalled LSP in that it is configured manually on each LSR; no signalling protocol is used. However, the structure of a static LSP is the same as a signalled LSP: an ingress LER adds a label and forwards the packet; one or more transit LSRs receive the packet, swap the label, and forward the packet; the packet exits the LSP at the egress LER. Unlike signalled LSPs, static LSPs do not perform penultimate hop popping. In a static LSP, the egress LER pops the label.

To establish a static LSR, you configure the ingress, transit, and egress LERs, manually specifying the labels to be applied at each hop. Since packets can start flowing into an LSP at the instant the LSP is enabled on the ingress LER, Foundry recommends that you configure and enable the static LSP on the transit and egress LERs prior to enabling the LSP on the ingress LER.

Configuring the Ingress LER in a Static LSP

The ingress LER in an LSP classifies the packet, applies a label to the packet, and forwards the packet to a transit LSR or to the egress LER. When configuring the ingress LER in a static LSP, you specify which packets can be forwarded through the LSP, the label to be applied to the packets, and the outbound interface through which the packets are forwarded. You can perform the following tasks:

- Creating the static LSP
- Specifying the outbound interface and label

- Specifying an egress LER for the static LSP
- Setting aliases for the egress LER (optional)
- Setting a Class of Service (CoS) value for the LSP (optional)
- Allocating bandwidth to the outbound interface (optional)
- Setting a metric for a static LSP (optional)
- Enabling a static LSP
- Disabling a static LSP
- Assigning packets to the static LSP (optional)

Creating a Static LSP

To create a static LSP and enter the static LSP configuration level, enter commands such as the following:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)#
```

Syntax: [no] static-lsp <name>

Specifying the Outbound Interface and Label

In label switching, the ingress LER adds a label to a packet and forwards it out an outbound interface. For static LSPs, you manually specify the outbound interface as well as the label that is applied to the packet.

For example, to cause packets in the LSP to be sent out POS interface 3/11 with label 123:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# out-segment pos 2/1 out-label 123
```

Syntax: out-segment pos <interface> out-label <label>

Syntax: out-segment ethernet <interface> out-label <label> next-hop <address>

The <interface> is an MPLS-enabled interface on the device. The <label> can be a number from 16 – 499999. (Labels 0 – 15 are reserved by the IETF.) The Foundry device dynamically assigns labels in the range 1024 – 499999 to interfaces used for signalled LSPs. For static LSPs, Foundry recommends that you assign a label in the range 16 – 1023. You can still assign a label in the range 1024 – 499999 to a static LSP, but it may conflict with a label assigned to an existing signalled LSP. If there is such a conflict, an error message is displayed.

The **next-hop** parameter only applies to static LSPs whose outbound interface is an Ethernet interface.

Specifying the Egress LER

Like a signalled LSP, each static LSP requires one and only one egress LER. The egress LER is the router from which packets exit the MPLS domain from this LSP. After the LSP is successfully established, the address of the egress LER is installed as an internal host route, allowing the ingress LER to direct BGP next-hop traffic into the LSP.

To specify 64.100.1.1 as the address of the egress LER for static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# to 64.100.1.1
```

Syntax: to <ip address>

The egress LER is the only required parameter in an LSP. All other parameters are optional.

Creating Aliases for the Egress LER

Traffic whose BGP next hop is the egress LER in the LSP is always forwarded along the LSP. In addition, you can configure one or more aliases for the egress LER, so that traffic whose BGP next hop matches one of the aliases is also forwarded along the LSP. Creating aliases is useful in situations where a BGP next-hop router is not MPLS capable. Traffic for the alias can be sent along the LSP to the egress LER, and then to the non-MPLS router using Layer 3 forwarding.

For example, to install 1.1.1.1/16 as an alias for the egress LER in static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# install 1.1.1.1/16
```

Syntax: install <ip address/mask>

Traffic whose BGP next hop matches 1.1.0.0/16 is then forwarded along the LSP to the egress LER. When it arrives at the egress LER, it can be forwarded to the actual BGP next hop.

Setting a Class of Service Value for the LSP

The 3-bit EXP field in the MPLS header can be used to define a Class of Service (CoS) value for packets travelling through the LSP. The CoS value is used to specify a priority for MPLS packets.

There are two ways a CoS value can be applied to packets travelling through an LSP:

- Use the Type of Service (ToS) field in the IP header. This is the default behavior. Specifically, the Foundry device copies the first three bits in the packet's ToS field to the to the CoS (EXP) field in the MPLS header. The ToS value maps to one of the four priority queues on the Foundry device.
- Manually set a CoS value for the LSP. The CoS value that you set is applied to the CoS (EXP) field in the MPLS header of all packets entering this LSP. This way all packets travelling through an LSP can be treated with the same priority as they travel the MPLS domain. You can assign the LSP a CoS value from 0 to 7.

To assign a CoS value of 7 (highest priority) to all packets traveling through LSP static_tunnel:

```
NetIron(config-mpls)# lsp static_tunnel
NetIron(config-mpls-static-lsp)# cos 7
```

Syntax: cos <number>

The CoS value can be an integer from 0 – 7.

Allocating Bandwidth to the Outbound Interface in a Static LSP

You can specify the amount of bandwidth allocated to an LSP, including the maximum and average rate of packets that will be traveling through the LSP. Unlike signalled LSPs, where you can allocate bandwidth at the ingress LER and signal the allocations to the other LSRs, in a static LSP, you manually allocate bandwidth at each LSR. At the ingress LER, you can allocate bandwidth for the outbound interface.

To set the maximum rate of packets that can go through the outbound interface of the LSP (in Kbits/sec):

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng max-rate 20
```

Syntax: traffic-eng max-rate <rate>

To set the average rate of packets that can go through the outbound interface of the LSP (in Kbits/sec):

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng mean-rate 10
```

Syntax: traffic-eng mean-rate <rate>

To set the maximum size (in bytes) of the largest burst the outbound interface of the LSP can send at the maximum rate:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng max-burst 10
```

Syntax: traffic-eng max-burst <bytes>

Assigning a Metric to the LSP

You can assign a metric to the static LSP, which can be used by routing protocols to determine the relative preference of the LSP compared to other hop-by-hop routes. An LSP metric can also be used to determine preference among several LSPs towards a destination.

To assign a metric of 5 to static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# metric 5
```

Syntax: metric <number>

By default, all LSPs have a metric of 1. A lower metric is preferred over a higher one. If there are multiple LSPs to the same destination LSR, and they share the same metric, the traffic load is shared among them.

Enabling a Static LSP

After you have set up each LSR – ingress, transit, and egress – in the static LSP, you can enable the LSP. Enabling the LSP is the final step in configuring it. You should enable the static LSP on the transit and egress LERs prior to enabling the static LSP on the ingress LER.

To enable static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# enable
```

Syntax: enable

Disabling a Static LSP

Disabling an LSP de-activates it, but does not remove the LSP from the device's configuration. (To remove a static LSP from the device's configuration, use the **no static-lsp** command.) To make changes to an active LSP, first disable the LSP, modify parameters on the LSP, and then enable the LSP.

To disable LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# disable
```

Syntax: disable

Assigning Packets to the Static LSP

When a packet enters the MPLS domain, the ingress LER classifies it and determines to which LSP (if any) the packet should be assigned. By default, packets whose BGP next-hop address is the egress LER (or one of its aliases) are always assigned to the LSP. In addition, you can cause packets whose destination address matches a specified IP address prefix to be assigned to an LSP.

For example, to assign packets matching prefix 30.1.1.0/24 to LSP static_tunnel:

```
NetIron(config-mpls)# ip route 30.1.1.0/24 static-lsp static_tunnel
```

Syntax: ip route <dest-ip-addr>/<mask-bits> lsp | static-lsp <name>

Assigning packets to an LSP in this way is similar to installing a static route. Instead of specifying an Ethernet, POS, or VE interface as the next hop, you specify the name of a static LSP; packets matching the specified prefix are assigned to the LSP.

Configuring a Transit LSR in a Static LSP

The transit LSR in an LSP receives a labelled packet, swaps the label, and forwards the packet to another transit LSR or to the egress LER. When configuring a transit LSR in a static LSP, you specify the inbound interface and label that identifies packets belonging to the LSP, a new label to be applied to the packets, and the outbound interface through which the packets are forwarded.

To configure a transit LSR in a static LSP, you can perform the following tasks:

- Configuring label swapping
- Allocating bandwidth to the inbound and outbound interfaces (optional)
- Enabling the static LSP
- Disabling the static LSP

Configuring Label Swapping

The transit LSR in an LSP receives a labelled packet on a specified inbound interface, swaps the label, and forwards the packet out an outbound interface.

In a static LSP, you specify the inbound interface and the label the packet must have in order to be forwarded along the LSP. When the transit LSR receives a packet on this inbound interface, and the packet has the specified inbound label, the LSR looks up the interface-label combination in its MPLS forwarding table. The inbound interface-label combination maps to an outbound interface-label combination. Using this mapping information, the transit LSR swaps the inbound label with the outbound label and forwards the packet out the outbound interface.

In a static LSP, you manually set both the inbound interface-label combination and outbound interface-label combination.

For example, the following command specifies interface POS 2/1 as the inbound interface and label 123 as the inbound label for static LSP static_tunnel.

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# in-segment pos 2/1 in-label 123
```

Syntax: in-segment pos <interface> in-label <label>

Each MPLS-enabled interface is associated with a unique label space. In other words, incoming labeled packets with a given label value on one interface are distinct from incoming labeled packets with the same label value from another interface.

The following command specifies interface POS 3/1 as the outbound interface and label 456 as the outbound label for the static LSP.

```
NetIron(config-mpls-static-lsp)# out-segment pos 3/1 out-label 456
```

Syntax: out-segment pos <interface> out-label <label>

The <label> can be a number from 16 – 499999. (Labels 0 – 15 are reserved by the IETF.) The Foundry device dynamically assigns labels in the range 1024 – 499999 to interfaces used for signalled LSPs. For static LSPs, Foundry recommends that you assign a label in the range 16 – 1023. You can still assign a label in the range 1024 – 499999 to a static LSP, but it may conflict with a label assigned to an existing signalled LSP. If there is such a conflict, an error message is displayed.

After the static LSP is enabled, if a packet with label 123 comes into interface POS 2/1, the label is replaced with label 456, and the packet is forwarded out interface POS 3/1.

Allocating Bandwidth to the Inbound and Outbound Interfaces in a Static LSP

You can specify the amount of bandwidth allocated to the inbound and outbound interfaces, including the maximum and average rate of packets that will be traveling through the interface. Allocating bandwidth to an interface allows the LSRs to determine how much bandwidth may be consumed by the interface, as well as how much available bandwidth resources can be advertised via OSPF-TE LSAs.

Unlike signalled LSPs, where you can allocate bandwidth at the ingress LER and propagate the allocations to the other LSRs, in a static LSP, you manually allocate bandwidth at each LSR.

To set the maximum rate of packets that can go through the inbound and outbound interfaces of the LSP (in Kbits/sec):

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng max-rate 20
```

Syntax: traffic-eng max-rate <rate>

To set the average rate of packets that can go through the inbound and outbound interfaces of the LSR (in Kbits/sec):

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng mean-rate 10
```

Syntax: traffic-eng mean-rate <rate>

To set the maximum size (in bytes) of the largest burst the inbound and outbound interfaces of the LSP can send at the maximum rate:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# traffic-eng max-burst 10
```

Syntax: traffic-eng max-burst <bytes>

Enabling a Static LSP

After you have set up each LSR – ingress, transit, and egress – for the static LSP, you can enable it. Enabling the LSP is the final step in configuring it. You should enable the LSP on the transit and egress LERs prior to enabling it on the ingress LER.

To enable static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# enable
```

Syntax: enable

Disabling a Static LSP

Disabling an LSP de-activates it, but does not remove the LSP from the device's configuration. (To remove a static LSP from the device's configuration, use the **no static-lsp <name>** command.) To make changes to an active LSP, first disable the LSP, modify parameters on the LSP, and then enable the LSP.

To disable LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# disable
```

Syntax: disable

Configuring the Egress LER in a Static LSP

The egress LER in an LSP receives a labelled packet and removes the label. The packet thus exits the MPLS domain and is forwarded using standard Layer 3 routing protocols. When configuring an egress LER in a static LSP, you specify the inbound interface and label that identifies packets belonging to the LSP.

To configure an egress LER in a static LSP, you can perform the following tasks:

- Configuring label popping
- Enabling the static LSP
- Disabling the static LSP

Configuring Label Popping

The egress LER removes (pops) the label from an MPLS-encapsulated packet. You specify the inbound interface and inbound label for the MPLS packets. Since the egress LER is the end of the path, you do not specify an outbound interface and label.

For example, the following command specifies interface POS 3/1 as the inbound interface and label 456 as the inbound label for static LSP static_tunnel.

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# in-segment pos 3/1 in-label 456
```

Syntax: in-segment pos <interface> in-label <label>

Packets coming into interface POS 3/1 that have label 456 are identified as packets belonging to this LSP. In the MPLS forwarding table, there is no corresponding outbound interface-label pair for this inbound interface pair, so the label is removed (popped). From this point forward, the packet is routed using standard Layer 3 routing protocols.

Enabling a Static LSP

After you have set up each LSR – ingress, transit, and egress – for the static LSP, you can enable it. Enabling the LSP is the final step in configuring it.

To enable static LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# enable
```

Syntax: enable

Disabling a Static LSP

Disabling an LSP de-activates it, but does not remove the LSP from the device's configuration. (To remove a static LSP from the device's configuration, use the **no static-lsp <name>** command.) To make changes to an active LSP, first disable the LSP, modify parameters on the LSP, and then enable the LSP.

To disable LSP static_tunnel:

```
NetIron(config-mpls)# static-lsp static_tunnel
NetIron(config-mpls-static-lsp)# disable
```

Syntax: disable

Displaying MPLS and RSVP Information

You can display the following information about the MPLS configuration on the Foundry device:

- Information about MPLS-enabled interfaces on the device
- Statistics about the MPLS-enabled interfaces
- MPLS summary information
- Contents of the Traffic Engineering Database (TED)
- Status information about signalled LSPs configured on the device
- Status information about static LSPs of which the device is a component
- The label applied at each hop in an LSP
- Contents of the MPLS routing table
- RSVP information, including the status of RSVP-enabled interfaces, session information, and statistics
- Information about OSPF-TE LSAs

Displaying Information About MPLS-Enabled Interfaces

To display information about the interfaces on the device that have been enabled for MPLS:

```
NetIron# show mpls interface
e1/1
  Maximum BW: 1000000 kbps, maximum reservable BW: 1000000 kbps
  Admin group: 0x0000003a ( 1 3 4 5)
  Reservable BW [priority] kbps:
    [0] 1000000    [1] 1000000    [2] 1000000    [3] 1000000
    [4] 1000000    [5] 1000000    [6] 1000000    [7] 1000000
  Last sent reservable BW [priority] kbps:
    [0] 0    [1] 0    [2] 0    [3] 0
    [4] 0    [5] 0    [6] 0    [7] 0
```

Syntax: show mpls interface

For each MPLS-enabled interface on the device, the following information is displayed:

Table 3.1: Output from the show mpls interface command

This Field...	Displays...
Maximum BW:	The maximum outbound bandwidth that can be used on the interface. This TLV reflects the actual physical bandwidth of the interface (155M for OC-3, 622M for OC-12, or 2488M for OC-48).

Table 3.1: Output from the show mpls interface command (Continued)

This Field...	Displays...
maximum reservable BW:	The maximum bandwidth that can be reserved on the interface. By default, the Maximum Reservable Bandwidth is the same as the Maximum Bandwidth for the interface. You can optionally change the reservable bandwidth to an amount greater than or equal to the maximum available bandwidth of the interface with the **traffic-eng reservable-bw** command.
Admin group:	The administrative group(s) to which this interface belongs, set with the **admin-group** command.
Reservable BW [priority] kbps:	The amount of bandwidth not yet reserved on the interface. Eight octets are displayed, indicating the amount of unreserved bandwidth (in kbits per second) that can be reserved with a hold priority of 0 through 7. The value in each of the octets is less than or equal to the maximum reservable bandwidth.
Last sent reservable BW [priority] kbps:	The values in the Unreserved Bandwidth TLV sent in the most recent OSPF-TE LSA. If the device is not sending out OSPF-TE LSAs for the interface, the unreserved bandwidth value for each of the priorities is 0.

Displaying MPLS Statistics

Statistics about MPLS packets on the Foundry device are gathered on a per-interface basis, rather than on a per-LSP basis. For example:

```
NetIron# show mpls statistics

PORT p2/1 packet stats:

          Tunnel entry tx =      0 Tunnel entry drop =      0
           Tunnel exit rx =      0  Tunnel exit drop =      0
               XC inbound =      0   XC inbound drop =      0
              XC outbound =      0  XC outbound drop =      0
                   VLL tx =            0
                   VLL rx =            0

PORT p2/2 packet stats:

          Tunnel entry tx =      0 Tunnel entry drop =      0
           Tunnel exit rx =      0  Tunnel exit drop =      0
               XC inbound =      0   XC inbound drop =      0
              XC outbound =      0  XC outbound drop =      0
                   VLL tx =            0
                   VLL rx =            0
```

Syntax: show mpls statistics

For each MPLS-enabled interface on the device, the **show mpls statistics** command displays the following information:

Table 3.2: Output from the show mpls statistics command

This Field...	Displays...
Tunnel entry tx	The number of packets that entered an LSP at this LSR and were forwarded out this interface; that is, the number of times the device, serving as an ingress LER, pushed a label onto an IP packet and forwarded the labelled packet out this interface.
Tunnel entry drop	The number of packets that entered an LSP at this LSR, but were dropped before they could be forwarded out this interface.
Tunnel exit rx	For static LSPs, the number of labelled packets that exited an LSP at this LSR and were forwarded out this interface as IP packets. For signalled LSPs, this will be always be 0 because of penultimate hop popping. Signalled LSPs pop the label at the next to last LSR in the LSP, and forward the packet to the egress LER as an IP packet. Since there is no label on the packet when it reaches the egress LER, there is no way to determine which LSP (if any) the packet is exiting.
Tunnel exit drop	The number of packets that exited an LSP at this LSR, but were dropped before they could be forwarded out this interface as IP packets.
XC inbound	The number of labelled packets that were received on this interface, and whose labels were swapped at this LSR.
XC inbound drop	The number of labelled packets that were received on this interface, but were dropped.
XC outbound	The number of labelled packets whose labels were swapped at this LSR, and were forwarded out this interface.
XC outbound drop	The number of labelled packets whose labels were swapped at this LSR, but were dropped before they could be forwarded out this interface.
VLL tx	The number of MPLS packets containing a VC label that have been sent out an MPLS-enabled interface.
VLL rx	The number of packets MPLS packets containing a VC label that have been received on an MPLS-enabled interface.

To clear the MPLS statistics counters:

```
NetIron# clear mpls statistics
```

Syntax: clear mpls statistics

Displaying MPLS Summary Information

You can display a summary of MPLS information, including the number of configured paths, signalled LSPs, and static LSPs for which this device is the ingress LER.

For example:

```
NetIron# show mpls summary

Path:
        Paths configured    =     0

Signaled LSPs:
        LSPs configured     =     1
        LSPs enabled        =     1
        LSPs operational    =     0

Static LSPs:
        LSPs configured     =     0
        LSPs enabled/oper.  =     0
```

Syntax: show mpls summary

Displaying the Contents of the Traffic Engineering Database

An LSR's Traffic Engineering Database (TED) contains topology information about nodes in an MPLS domain and the links connecting them. This topology information is obtained from OSPF traffic engineering (OSPF-TE) LSAs. OSPF-TE LSAs have special extensions that contain information about an MPLS-enabled interface's traffic engineering metric, bandwidth reservations, and administrative group memberships.

An LSR, when configured to do so, floods OSPF-TE LSAs for its MPLS-enabled interfaces to its neighboring routers in the OSPF area. Other LSRs store the information from the OSPF-TE LSAs in their own Traffic Engineering Databases, allowing each LSR in the OSPF area to maintain an identical TED describing the MPLS topology. The topology information in the TED is used by the CSPF process when it calculates traffic-engineered paths for signalled LSPs.

You can display the contents of an LSR's TED. For example:

```
NetIron# show mpls ted data
AreaId: 0
  NodeID: 2.2.2.2, Type: Router
    Type: M/A, To: 10.1.1.3, Local: 10.1.1.2, Remote: 0.0.0.0
  NodeID: 3.3.3.3, Type: Router
    Type: P2P, To: 6.6.6.6, Local: 40.1.1.1, Remote: 40.1.1.2
    Type: M/A, To: 10.1.1.3, Local: 10.1.1.3, Remote: 0.0.0.0
    Type: M/A, To: 20.1.1.2, Local: 20.1.1.1, Remote: 0.0.0.0
  NodeID: 10.1.1.3, Type: Network
    Type: M/A, To: 1.1.1.1, Local: 0.0.0.0, Remote: 0.0.0.0
    Type: M/A, To: 2.2.2.2, Local: 0.0.0.0, Remote: 0.0.0.0
    Type: M/A, To: 3.3.3.3, Local: 0.0.0.0, Remote: 0.0.0.0
  NodeID: 30.1.1.2, Type: Network
    Type: M/A, To: 1.1.1.1, Local: 0.0.0.0, Remote: 0.0.0.0
    Type: M/A, To: 6.6.6.6, Local: 0.0.0.0, Remote: 0.0.0.0
```

Syntax: show mpls ted data

The following table describes the output of the **show mpls ted data** command.

Table 3.3: Output from the show mpls ted data command

This Field...	Displays...
AreaId:	The ID of this OSPF area.
NodeID:	The ID of the node. For Router nodes, can be any interface address or a loopback interface address on the LSR. For Network nodes, this is the router ID of the network's designated router.
[node] Type:	The type of node. The node type can be either Router or Network: Router Indicates the node is an actual LSR. Network Indicates the node represents a multi-access network.
[link] Type:	The type of link. The link type can be either P2P or M/A: P2P Indicates this is a point-to-point link. M/A Indicates the link is a broadcast, multi-access network.
To:	The ID of the node at the end of this link.
Local:	The address of the interface used to reach the remote node. For M/A link types, this is always 0.0.0.0.
Remote:	The address of the interface on the remote node that is connected to the local node. For M/A link types, this is always 0.0.0.0.

To display more detailed information about each node in the TED:

```
NetIron# show mpls ted data detail
AreaId: 0
  NodeID: 2.2.2.2, Type: Router
    Type: M/A, To: 10.1.1.3, Local: 10.1.1.2, Remote: 0.0.0.0
      Color: 0x00000007
      Metric: 1
      Max BW: 155000 kbps
      Reservable BW: 155000 kbps
      Available BW [priority] kbps:
        [0] 155000       [1] 155000       [2] 155000       [3] 155000
        [4] 155000       [5] 155000       [6] 155000       [7] 155000
  NodeID: 1.1.1.1, Type: Router
    Type: M/A, To: 10.1.1.3, Local: 10.1.1.1, Remote: 0.0.0.0
      Color: 0x00000007
      Metric: 1
      Max BW: 155000 kbps
      Reservable BW: 155000 kbps
      Available BW [priority] kbps:
        [0] 155000       [1] 155000       [2] 155000       [3] 155000
        [4] 155000       [5] 155000       [6] 155000       [7] 155000
    Type: M/A, To: 30.1.1.2, Local: 30.1.1.1, Remote: 0.0.0.0
      Color: 0x00000007
      Metric: 1
      Max BW: 155000 kbps
      Reservable BW: 155000 kbps
      Available BW [priority] kbps:
        [0] 155000       [1] 155000       [2] 155000       [3] 155000
        [4] 155000       [5] 155000       [6] 155000       [7] 155000
```

Syntax: show mpls ted data detail

In addition to the information described in Table 3.3, the **show mpls ted data detail** command displays the following:

Table 3.4: Output from the show mpls ted data detail command

This Field...	Displays...
Color:	The administrative group(s) to which this interface belongs.
Metric:	The traffic engineering metric for the interface (by default, this is equal to the OSPF link cost).
Max BW:	The maximum outbound bandwidth that can be used on the interface. This is the actual physical bandwidth of the interface (155M for OC-3, 622M for OC-12, or 2488M for OC-48).
Reservable BW:	The maximum bandwidth that can be reserved on the interface. By default, the Maximum Reservable Bandwidth is the same as the Maximum Bandwidth for the interface.

Table 3.4: Output from the show mpls ted data detail command (Continued)

This Field...	Displays...
Available BW [priority] kbps:	The amount of bandwidth not yet reserved on the interface. Eight octets are displayed, indicating the amount of unreserved bandwidth (in kbits per second) that can be reserved with a hold priority of 0 through 7. The value in each of the octets is less than or equal to the maximum reservable bandwidth.

Displaying Signalled LSP Status Information

You can display status information about signalled LSPs for which the device is the ingress LER. For example:

```
NetIron# show mpls lsp
*: The LSP is taking a Secondary Path
                               Admin Oper  Tunnel   Up/Dn Retry Active
Name           To              State State Intf     Times No.   Path
t1             3.3.3.3         UP    UP*   tnl1     1     5     v2
```

Syntax: show mpls lsp [brief]

NOTE: The **show mpls lsp brief** command displays the same information as the **show mpls lsp** command.

The following table describes the output of the **show mpls lsp** command.

Table 3.5: Output from the show mpls lsp command

This Field...	Displays...
Name	The name of the LSP. LSPs are displayed in alphabetical order.
To	The egress LER for the LSP.
Admin State	The administrative state of the LSP. Once you activate the LSP with the **enable** command, the administrative state changes from DOWN to UP.
Oper State	The operational state of the LSP. This field indicates whether the LSP has been established through signalling and is capable of having packets forwarded through it. There may be a short period of time after you enable the LSP that the administrative state of the LSP is UP, but the operational state is DOWN. Once the LSP has been established through signalling, both the administrative state and the operational state will be UP.
Tunnel Intf	The MPLS tunnel interface port ID.
Up/Dn Times	The number of times the operational state of the LSP's primary path has transitioned from DOWN to UP.
Retry No.	The number of attempts the ingress LER has made to connect to the egress LER.
Active Path	The path currently in use for this LSP. Dashes (--) indicate that there is no named path for the LSP or that the LSP has not yet been established over the named path.

To display more detailed information about the status of the LSPs for which the device is the ingress LER:

```
NetIron# show mpls lsp detail
LSP t1
  To 3.3.3.3, admin: UP, status: UP, tunnel interface: tnl1
  Times primary LSP goes up: 1
  Metric: 1, number of installed aliases: 0
  Maximum retries: 0, no. of retries: 3
  Pri. path: dir, active: no
  Sec. path: v2, active: yes
    Hot-standby: no, status: up
  Setup priority: 7, hold priority: 0
  Max rate: 0 kbps, mean rate: 0 kbps, max burst: 0 bytes
  Constraint-based routing enabled: yes
  Tie breaking: random, hop limit: 0
  Explicit path hop counts: 2
    10.10.10.2 (S) -> 20.20.20.2 (S)
  Recorded routes:
    10.10.10.2 -> 20.20.20.2
```

Syntax: show mpls lsp detail | <name>

NOTE: The **show mpls lsp <name>** command displays detailed information about a specific LSP.

Table 3.6: Output from the show mpls lsp detail command

This Field...	Displays...
Name	The name of the LSP. LSPs are displayed in alphabetical order.
To	The egress LER for the LSP.
admin:	The administrative state of the LSP. Once you activate the LSP with the **enable** command, the administrative state changes from DOWN to UP.
status:	The operational state of the LSP. This field indicates whether the LSP has been established through signalling and is capable of having packets forwarded through it. If the status of the LSP is DOWN, the reason why the LSP is down is shown in parentheses. There may be a short period of time after you enable the LSP that the administrative state of the LSP is UP, but the status is DOWN. Once the LSP has been established through signalling, both the administrative state and the status will be UP.
Times primary LSP goes up:	The number of times the status of the LSP's primary path has transitioned from DOWN to UP.
Metric	The metric for the LSP, configured with the **metric** command.
no. of installed aliases:	The number of aliases that have been installed for the egress LER.
Max retries:	The maximum number of attempts the ingress LER will attempt to connect to the egress LER, set with the **retry-limit** command.
no. of retries:	The number of attempts the ingress LER has made to connect to the egress LER.
Pri. path:	The name of the primary path for this LSP and whether the path is currently active.
Sec. path:	The name of the secondary path for this LSP and whether the path is currently active.

Table 3.6: Output from the show mpls lsp detail command (Continued)

This Field...	Displays...
Hot-standby:	Whether the secondary path is a hot-standby path.
status:	The operational state of the secondary path.
Setup priority	The configured setup priority for the LSP.
hold priority	The configured hold priority for the LSP.
Max rate:	The maximum rate of packets that can go through the LSP (in Kbits/sec), set with the **traffic-eng max-rate** command.
mean rate:	The average rate of packets that can go through the LSP (in Kbits/sec), set with the **traffic-eng mean-rate** command.
max burst:	The maximum size (in bytes) of the largest burst the LSP can send at the maximum rate, set with the **traffic-eng max-burst** command.
Constraint-based routing enabled:	Whether CSPF is in effect for the LSP.
Tie breaking:	The tie-breaking method CSPF uses to select a path from a group of equal-cost paths to the egress LER, set with the **tie-breaking** command.
hop limit:	The maximum number of hops a path calculated by CSPF can have, set with the **hop-limit** command.
Explicit path hop counts:	The number of explicit hops configured for the LSP, the addresses of the hops, and whether the hops are strict (S) or loose (L).
Recorded routes:	The addresses recorded by the RECORD_ROUTE object during RSVP signalling.

NOTE: The part of the output starting from the Setup priority line is relevant to the currently active path. The settings for the secondary path may be different from the primary path. If no secondary path is configured for this LSP, the output always describes the primary path.

Displaying Static LSP Status Information

You can display status information about static LSPs of which the Foundry device is a component. For example:

```
NetIron# show mpls static-lsp
                       Tunnel In     In     Out    Out
Name            State  Intf   Port   Label  Port   Label  To
1               UP     tnl2   --     --     p2/1   200    2.2.2.2
2               UP     --     p1/1   100    p1/2   500
```

Syntax: show mpls static-lsp [brief]

NOTE: The **show mpls static-lsp brief** command displays the same information as the **show mpls static-lsp** command.

In this example, the Foundry device is a component of four static LSPs: it is the ingress LER for static LSP sl3, a transit LSR for static LSP sl2, and the egress LER for static LSPs sl1 and sl4.

The following table describes the output of the **show mpls static-lsp** command.

Table 3.7: Output from the show mpls static-lsp command

This Field...	Displays...
Name	The name of the static LSP. LSPs are displayed in alphabetical order.
State	The operational state of the static LSP. Once the static LSP has been activated on all its LSRs with the **enable** command, the State for the LSP is UP. If the static LSP has not been activated on all its LSRs, the State for the LSP is DOWN.
Tunnel Intf	The MPLS tunnel interface port ID. A value appears in this field only if this device is the ingress LER for the static LSP.
In Port	The inbound interface for the static LSP, set with the **in-segment** parameter. A value is shown in this field only if this device is a transit LSR or the egress LER for the static LSP.
In Label	The inbound label for packets in the static LSP, set with the **in-label** parameter. A value is shown in this field only if this device is a transit LSR or the egress LER for the static LSP.
Out Port	The outbound interface for the static LSP, set with the **out-segment** parameter. A value is shown in this field only if this device is the ingress LER or a transit LSR for the static LSP.
Out Label	The outbound label applied to packets in the static LSP, set with the **out-label** parameter. A value is shown in this field only if this device is the ingress LER or a transit LSR for the static LSP.
To	The egress LER for the static LSP. A value is shown in this field only if this device is the ingress LER for the static LSP.

To display more detailed information about static LSPs of which the Foundry device is a component:

```
NetIron# show mpls static-lsp detail
Static LSP 1
  Admin/oper: UP
  Placement: INGRESS, To: 2.2.2.2
  Out-segment: port p2/1, label 200
  Metric: 1, number of installed aliases: 0
  Max rate: 1000 kbps, mean rate: 100 kbps, max burst: 0
Static LSP 2
  Admin/oper: UP
  Placement: TRANSIT
  In-segment: port p1/1, label 100
  Out-segment: port p1/2, label 500
  Metric: 1, number of installed aliases: 0
  Max rate: 250 kbps, mean rate: 250 kbps, max burst: 0
```

Syntax: show mpls static-lsp detail | <name>

NOTE: The **show mpls static-lsp <name>** command displays detailed information about a specific static LSP.

The following table describes the output of the **show mpls static-lsp** command. The static LSPs are displayed in alphabetical order.

Table 3.8: Output from the show mpls static-lsp detail command

This Field...	Displays...
Admin/oper:	The administrative state of the LSP. Once you activate the LSP with the **enable** command, the administrative state changes from DOWN to UP.
Placement:	The role the device has in the static LSP. This can be INGRESS, TRANSIT, or EGRESS.
In-segment:	For transit LSRs and egress LERs, the inbound interface and the label the packet must have in order to be forwarded along the LSP, set with the **in-segment** and **in-label** parameters.
Out-segment:	For ingress LERs and transit LSRs, the outbound interface and the label applied to the packet, set with the **out-segment** and **out-label** parameters.
Metric:	The metric for the LSP, configured with the **metric** command.
no. of installed aliases:	The number of aliases that have been installed for the egress LER.
Max rate:	The maximum rate of packets that can go through the LSP (in Kbits/sec), set with the **traffic-eng max-rate** command.
mean rate:	The average rate of packets that can go through the LSP (in Kbits/sec), set with the **traffic-eng mean-rate** command.
max burst:	The maximum size (in bytes) of the largest burst the LSP can send at the maximum rate, set with the **traffic-eng max-burst** command.

Displaying Label Information

The **traceroute** command has been enhanced to display label information for each hop in an LSP. For example:

```
NetIron# traceroute 10.1.1.1
1   <1ms   <1ms   <1ms   20.1.1.1   [MPLS: Label 1024, CoS 0, TTL 1]
2   <1ms   <1ms   <1ms   30.1.1.1   [MPLS: Label 2048, CoS 0, TTL 1]
3   <1ms   <1ms   <1ms   40.1.1.1   [MPLS: Label 3000, CoS 0, TTL 1]
4   <1ms   <1ms   <1ms   10.1.1.1
```

Syntax: traceroute <host-ip-addr> [maxttl <value>] [minttl <value>] [numeric] [timeout <value>] [source-ip <ip addr>]

NOTE: If you suppress TTL propagation for the LSP (by entering the **no propagate-ttl** command on both the ingress LER and the LSR that pops the label) the LSP appears as a single hop.

Displaying the MPLS Routing Table

The MPLS routing table is used to store routes to egress LERs. When BGP performs next-hop resolution, it first looks into the MPLS routing table. If it cannot find a next hop in the MPLS routing table, it looks in the main routing table.

To display the contents of the MPLS routing table:

```
NetIron# show mpls route
Total number of MPLS tunnel routes: 2
Start index: 1
      Destination       NetMask           Gateway           Port    Cost
      20.1.1.2          255.255.255.255   20.1.1.2          tnl0       1
      30.1.1.2          255.255.255.255   20.1.1.2          tnl0       1
```

Syntax: show mpls route

Table 3.9: Output from the show mpls route command

This Field...	Displays...
Destination	The destination for the route. This can be either the address of the egress LER in an LSP, or a configured alias.
NetMask	The network mask for the route. If the destination address is the egress LER, this is a 32-bit mask. If the destination address is an alias, it is the network mask configured with the **install** command.
Gateway	The address of the egress LER in the LSP. If the destination address is not a network alias, the gateway is the same as the destination address.
Port	The MPLS tunnel interface associated with the LSP.
Cost	The metric for the LSP, set with the **metric** command in the LSP's configuration.

Displaying RSVP Information

You can display RSVP version information, the status of RSVP interfaces, RSVP session information, and RSVP statistics.

Displaying the RSVP Version

To display the RSVP version number, as well as the refresh interval and refresh multiple:

```
NetIron# show mpls rsvp
Resource ReSerVation Protocol, version 1. rfc2205
RSVP protocol          = Enabled
R (refresh interval) = 30 seconds
K (refresh multiple) = 3
```

Syntax: show mpls rsvp

Displaying the Status of RSVP interfaces

To display the status of RSVP on devices where it is enabled:

```
NetIron# show mpls rsvp interface
Interface State
     p2/1 Up
     p2/2 Dn
     p4/1 Dn
     p4/2 Dn
```

Syntax: show mpls rsvp interface [brief]

NOTE: The **show mpls rsvp interface brief** command displays the same information as the **show mpls rsvp interface** command.

In this example, interfaces POS 2/1, 2/2, 4/1, and 4/2 have been enabled for RSVP. Of these interfaces, interface POS 2/1 can actively send and receive RSVP messages.

To display detailed information about RSVP-enabled interfaces:

```
NetIron# show mpls rsvp interface detail
 Interface State
      P2/1 Up
   PacketType          Total                  Since last clear
                 Sent      Received         Sent      Received
   Path          8288             0            1             0
   PathErr          0             0            0             0
   PathTear         0             0            0             0
   Resv             0          3372            0             0
   ResvErr       1685             0            0             0
   ResvTear         0             0            0             0
```

Syntax: show mpls rsvp interface detail

For each RSVP-enabled interface, the following information is displayed:

Table 3.10: Output from the show mpls rsvp interface detail command

This Field...	Displays...
Path	The number of Path messages sent and received on the interface. Path messages store information about the state of the path along the LSRs in the LSP.
PathErr	The number of PathErr messages sent and received on the interface.
PathTear	The number of PathTear messages sent and received on the interface. PathTear messages cause path states to be deleted.
Resv	The number of Resv messages sent and received on the interface. Resv messages include FF (Fixed Filter), WF (Wildcard Filter), and SE (Shared Explicit) messages.
ResvErr	The number of ResvErr messages sent and received on the interface.
ResvTear	The number of ResvTear messages sent and received on the interface. ResvTear messages cause reservation states to be deleted.

To clear the RSVP statistics counters:

```
NetIron# clear mpls rsvp statistics
```

Syntax: clear mpls rsvp statistics

This command resets the counters listed under “Since last clear” for the **show mpls rsvp interface detail** and **show mpls rsvp statistics** commands.

Displaying RSVP Session Information

To display RSVP session information:

```
NetIron# show mpls rsvp session
To               From             State Style Labelin Labelout LSPname
Ingress RSVP:    1 session(s)
20.1.1.2         10.1.1.1         Up    FF    -       1025     t1


Transit RSVP:    0 session(s)

Egress RSVP:     1 session(s)
10.1.1.1         20.1.1.2         Up    FF    1024    -        t2
```

Syntax: show mpls rsvp session [brief | detail]

NOTE: The **show mpls rsvp session brief** command displays the same information as the **show mpls rsvp session** command.

The following table describes the output of the **show mpls rsvp session** command.

Table 3.11: Output from the show mpls rsvp session command

This Field...	Displays...
Ingress RSVP:	Information about ingress RSVP sessions.
Transit RSVP:	Information about transit RSVP sessions.
Egress RSVP:	Information about egress RSVP sessions.
To	Destination (egress LER) of the session.
From	Source (ingress LER) of the session.
Style	The RSVP reservation style. Possible values are FF (Fixed Filter), WF (Wildcard Filter), or SE (Shared Explicit).
Labelin	The label for inbound packets on this LSP.
Labelout	The label applied to outbound packets on this LSP.
LSPname	The name of the LSP.

To display more detailed information about RSVP sessions:

```
NetIron# show mpls rsvp session detail
Ingress RSVP:    1 session(s)
To               From             State Style Labelin Labelout LSPname
20.1.1.2         10.1.1.1         Up    FF    -       1025     t1
  Time left in seconds (PATH refresh: 10  RESV refresh: 6)
  Tspec: peak 0 kbps rate 0 kbps size 0 bytes m 20 M 65535
  PATH sentto:  10.1.1.2        (p2/1          )
  RESV rcvfrom: 10.1.1.2        (p2/1          )


Transit RSVP:    0 session(s)

Egress RSVP:     1 session(s)
To               From             State Style Labelin Labelout LSPname
10.1.1.1         20.1.1.2         Up    FF    1024    -        t2
  Time left in seconds (PATH refresh: 12  RESV refresh: 36)
  Tspec: peak 0 kbps rate 0 kbps size 0 bytes m 20 M 4470
  PATH rcvfrom: 10.1.1.2        (p2/1          )
```

Syntax: show mpls rsvp session detail

The show mpls rsvp session detail displays the same information described in Table 3.11, as well as the following:

Table 3.12: Output from the show mpls rsvp session detail command

This Field...	Displays...
Time left in seconds:	The amount of time left for the PATH or RESV refreshes.
Tspec:	Traffic engineering specification for the LSP, including the max-rate, mean rate, and burst rate settings.
PATH sentto:	Address of the next LSR in the LSP, and the interface used to reach this LSR.
PATH rcvfrom:	Address of the previous LSR in the LSP, and the interface used to reach this LSR.

Displaying RSVP Statistics

The Foundry device constantly gathers RSVP statistics. RSVP statistics are collected from the time RSVP is enabled, as well as from the last time the RSVP statistics counters were cleared.

To display RSVP statistics:

```
NetIron# show mpls rsvp statistics
                       Total                  Since last clear
 PacketType        Sent      Received        Sent      Received
  Path             4         4               4         4
  Resv             4         4               4         4
  PathErr          0         0               0         0
  ResvErr          0         0               0         0
  PathTear         0         0               0         0
  ResvTear         0         0               0         0
  ResvConf         0         0               0         0

 Errors                 Total                 Since last clear
  Rcv pkt bad length    0                     0
  Rcv pkt unknown type  0                     0
  Rcv pkt bad version   0                     0
  Rcv pkt bad cksum     0                     0
  Memory alloc fail     0                     0
```

Syntax: show mpls rsvp statistics

The following table describes the output of the **show mpls rsvp statistics** command.

Table 3.13: Output from the show mpls rsvp statistics command

This Field...	Displays...
Path	The number of PATH messages sent and received. PATH messages store information about the state of the path along the LSRs in the LSP.
Resv	The number of RESV messages sent and received. RESV messages include FF (Fixed Filter), WF (Wildcard Filter), and SE (Shared Explicit) messages.
PathErr	The number of PathErr messages sent and received.
ResvErr	The number of ResvErr messages sent and received.
PathTear	The number of PathTear messages sent and received. PathTear messages cause path states to be deleted.
ResvTear	The number of ResvTear messages sent and received. ResvTear messages cause reservation states to be deleted.
ResvConf	The number of reservation confirmation messages sent and received.
Rcv pkt bad length	The number of times a packet was not processed because it was the wrong length.
Rcv pkt unknown type	The number of times an RSVP packet was not processed because it was not one of the types defined in RFC 2205.
Rcv pkt bad version	The number of times a packet was not processed because it was an RSVP version other than 1
Rcv pkt bad cksum	The number of times a packet was not processed because of a bad RSVP checksum.
Memory alloc fail	The number of times a packet was not processed because RSVP memory allocation failed on the Foundry device.

To clear the RSVP statistics counters:

```
NetIron# clear mpls rsvp statistics
```

Syntax: clear mpls rsvp statistics

This command resets the counters listed under “Since last clear” for the **show mpls rsvp interface detail** and **show mpls rsvp statistics** commands.

Displaying Information About OSPF-TE LSAs

To display information about OSPF-TE LSAs:

```
NetIron# show ip ospf database link-state opaque-area

Area ID         Type    LS ID           Adv Rtr         Seq(Hex)  Age    Cksum
0               OpAr    1.0.0.0         3.3.3.3         80000006  1337  0x1a19
  Area-opaque TE LSA
  1 - router address (len 4): 3.3.3.3

Area ID      Type    LS ID           Adv Rtr         Seq(Hex)   Age      Cksum
0            OpAr    1.0.0.2         2.2.2.2         80000007  1333     0x88f1
  Area-opaque TE LSA
  2 - link (len 100):
    1 - link type (len 1): point-to-point(1)
    2 - link ID (len 4): 1.1.1.1
    3 - local i/f ip addr (len 4): 10.1.1.2
    4 - remote i/f ip addr (len 4): 10.1.1.1
    5 - TE metric (len 4):
    6 - max BW (len 4): 2372 Mbits/sec
    7 - max reservable BW (len 4): 2372 Mbits/sec
    8 - unreserved BW (len 32):
      Priority 0: 2372 Mbits/sec
      Priority 1: 2372 Mbits/sec
      Priority 2: 2372 Mbits/sec
      Priority 3: 2372 Mbits/sec
      Priority 4: 2372 Mbits/sec
      Priority 5: 2372 Mbits/sec
      Priority 6: 2372 Mbits/sec
      Priority 7: 2372 Mbits/sec
    9 - color (len 4): 0
```

Syntax: show ip ospf database link-state opaque-area

Displaying Information About IGP Shortcuts

To display routes that are using an LSP as an IGP shortcut:

```
NetIron# show ip route tunnel
Start index: 1  B:BGP D:Connected  R:RIP  S:Static  O:OSPF *:Candidate default
      Destination       NetMask           Gateway           Port   Cost   Type
2     2.2.2.2           255.255.255.255   2.2.2.2           tnl0   1       O
                                          LSP name: tunnel1 Signaling: RSVP
```

Syntax: show ip route tunnel

MPLS Sample Configurations

This section presents examples of typical MPLS configurations. The following sample configurations are presented:

- "LSP Between BGP Neighbors in an AS" on page 3-58
- "LSP with an Alias for a Non-MPLS Destination" on page 3-61
- "LSP with Redundant Paths" on page 3-63
- "Static LSP Configuration" on page 3-66

LSP Between BGP Neighbors in an AS

Figure 3.12 depicts a signalled LSP between BGP neighbors in an autonomous system. The LSP serves as a shortcut for BGP next-hop traffic within the AS.

Figure 3.12   LSP between two BGP neighbors

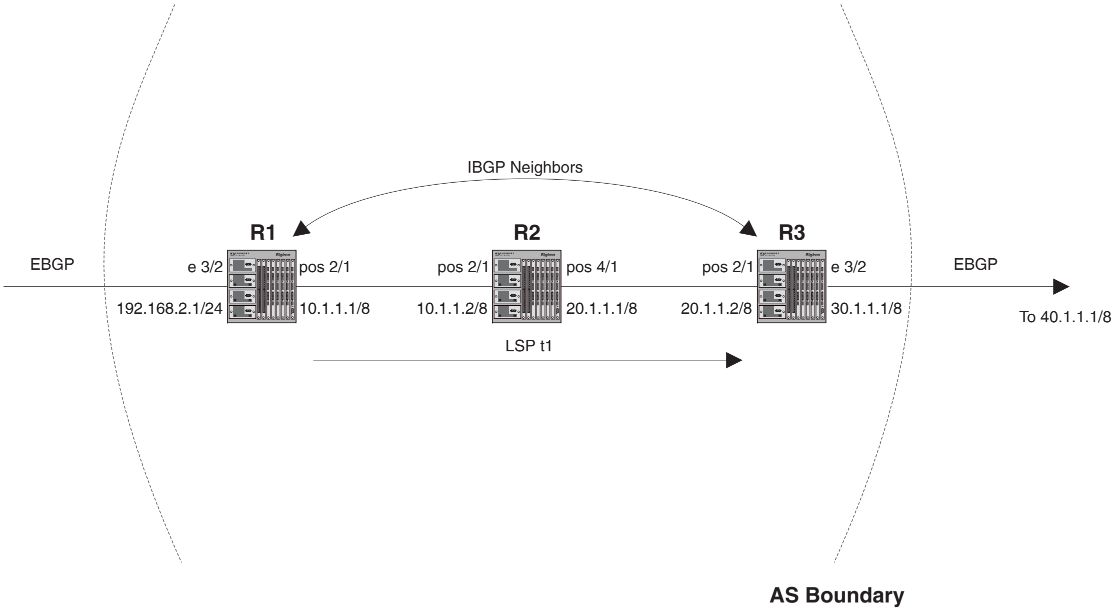

Router R3 advertises to its BGP neighbor R1 that it can reach network 40.1.1.1/8. To send R3 packets destined for this network, R1 needs to determine how to reach R3. Without MPLS, the route to R3 is determined using the IGP shortest path. With MPLS, you can create an LSP that specifies an alternate path between R1 and R3.

In this example, R1 is configured with a signalled LSP that specifies R3 as the destination. When this LSP is enabled, a path is established through R2 to reach R3. Packets whose BGP next hop is R3 are assigned to this LSP. To do this, a route to R3 is placed in R1's MPLS routing table.

In the MPLS routing table, a destination address is associated with an LSP to be used to reach that destination. BGP uses information in the MPLS routing table and the IP routing table (inet.0) to determine where to forward a packet. After determining the next hop address for a packet, BGP looks in the MPLS routing table for a match for this address. If a match is found, the packet is assigned to the LSP associated with the address. If a match is not found, BGP uses information in the IP routing table, which is generally the IGP shortest path. A route in the MPLS routing table is always favored over the IGP shortest path in the IP routing table.

If BGP finds a route for the next hop address in the MPLS routing table, it places the route into the main routing table. In this example, when BGP resolves a next hop address as R3, BGP looks in the MPLS routing table for a route to R3. In R1's MPLS routing table, a route to R3 is associated with LSP t1. BGP places this route into the main routing table. LSP t1 is installed as the outgoing "interface" for address prefixes for which R3 is the next hop.

Consequently, packets going to 40.1.1.1/8 (as well as any other prefixes for which R3 is the BGP next hop) are always assigned to LSP t1.

You can view information in the MPLS routing table with the **show mpls route** command. The following is an example of this command entered on R1.

```
NetIron# show mpls route
Total number of MPLS tunnel routes: 1
Start index: 1
      Destination       NetMask           Gateway           Port   Cost
      20.1.1.2          255.255.255.255   20.1.1.2          t1      0
```

The output of this command indicates that destination address 20.1.1.2, an interface on R3, is associated with LSP t1.

Router R1

The following commands configure Router R1 in Figure 3.12:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# lsp t1
NetIron(config-mpls-lsp)# to 20.1.1.2
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.1 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 192.168.2.1 255.255.255.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router bgp
NetIron(config-bgp-router)# local-as 100
NetIron(config-bgp-router)# neighbor 192.168.2.2 remote-as 50
NetIron(config-bgp-router)# neighbor 20.1.1.2 remote-as 100
NetIron(config-bgp-router)# network 192.168.2.0 255.255.255.0
NetIron(config-bgp-router)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R2

The following commands configure Router R2 in Figure 3.12:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface pos 4/1
NetIron(config-posif-4/1)# ip address 20.1.1.1 255.0.0.0
NetIron(config-posif-4/1)# ip ospf area 1
NetIron(config-posif-4/1)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R3

The following commands configure Router R3 in Figure 3.12:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# lsp t2
NetIron(config-mpls-lsp)# to 10.1.1.1
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 20.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 30.1.1.1 255.0.0.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router bgp
NetIron(config-bgp-router)# local-as 100
NetIron(config-bgp-router)# neighbor 10.1.1.1 remote-as 100
NetIron(config-bgp-router)# neighbor 40.1.1.2 remote-as 200
NetIron(config-bgp-router)# network 40.0.0.0 255.0.0.0
NetIron(config-bgp-router)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

LSP with an Alias for a Non-MPLS Destination

If a BGP next hop router is not MPLS capable, you can configure an LSP that treats the egress LER as an alias for the BGP next hop router. Traffic whose BGP next hop matches the alias is forwarded along the LSP to the egress LER. Once this traffic reaches the egress LER, it can be forwarded to the BGP next hop router using the IGP shortest path. Figure 3.13 shows an example of this kind of configuration.

Figure 3.13 LSP with an alias for a non-MPLS BGP next hop router

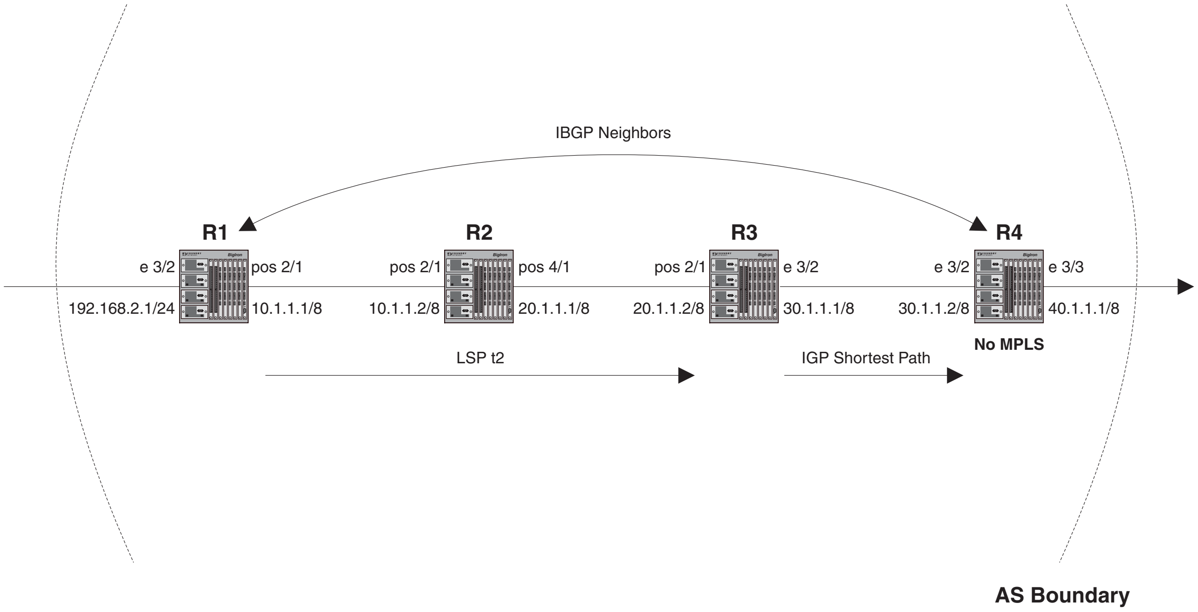

In this example, R1, R2, and R3 are MPLS capable, but R4 is not. An LSP is configured on R1 that uses a path through R2 to reach R3. In the LSP's configuration, R4 is specified as an alias of R3. This means that when R1 receives packets bound for R4, it assigns them to the LSP, sending them to R3. When the packets exit the LSP at R3, they can be forwarded to R4 using standard hop-by-hop routing.

When the LSP is enabled, routes to both R3 and R4 are placed in R1's MPLS routing table, specifying that this LSP be used for packets going to either of these destinations. After BGP determines that R4 is the next hop address for a packet, it looks in the MPLS routing table for a match for this address. Finding a match, BGP places the route into the main routing table.

In this example, BGP installs LSP t2 in R1's main routing table as the outgoing "interface" for address prefixes for which R4 is the next hop. Consequently, packets going from R1 to 40.1.1.1/8 (as well as any other prefixes for which R4 is the BGP next hop) are always assigned to LSP t2. When these packets reach R3 – LSP t2's egress LER – they are forwarded to R4 using the IGP shortest path.

The **show mpls route** command identifies which destinations in the MPLS routing table are aliases. For example, the following is the contents of R1's MPLS routing table:

```
NetIron# show mpls route
Total number of MPLS tunnel routes: 2
Start index: 1
      Destination       NetMask           Gateway           Port    Cost
      20.1.1.2          255.255.255.255   20.1.1.2          t2       0
      40.1.1.1          255.0.0.0         20.1.1.2          t2       0
```

Address 20.1.1.2/32 is an interface on R3, and address 40.1.1.1/8 is an interface on R4. Both addresses in the Destination column are associated with LSP t2; packets going to either destination are assigned to this LSP. The address in the Gateway column indicates the actual egress LER for the LSP. When the Gateway address is different from the Destination address, it means that the Destination address is an alias.

Router R1

The following commands configure Router R1 in Figure 3.13:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# lsp t2
NetIron(config-mpls-lsp)# to 20.1.1.2
NetIron(config-mpls-lsp)# install 40.1.1.1/8
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.1 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 192.168.2.1 255.255.255.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router bgp
NetIron(config-bgp-router)# local-as 100
NetIron(config-bgp-router)# neighbor 192.168.2.2 remote-as 50
NetIron(config-bgp-router)# neighbor 30.1.1.2 remote-as 100
NetIron(config-bgp-router)# network 192.168.2.0 255.255.255.0
NetIron(config-bgp-router)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R2

The following commands configure Router R2 in Figure 3.13:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface pos 4/1
NetIron(config-posif-4/1)# ip address 20.1.1.1 255.0.0.0
NetIron(config-posif-4/1)# ip ospf area 1
NetIron(config-posif-4/1)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R3

The following commands configure Router R3 in Figure 3.13:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 20.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 30.1.1.1 255.0.0.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R4

The following commands configure Router R4 in Figure 3.13:

```
NetIron(config)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 30.1.1.2 255.0.0.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# interface e 3/3
NetIron(config-if-e100-3/3)# ip address 40.1.1.1 255.0.0.0
NetIron(config-if-e100-3/3)# exit

NetIron(config)# router bgp
NetIron(config-bgp-router)# local-as 100
NetIron(config-bgp-router)# neighbor 10.1.1.1 remote-as 100
NetIron(config-bgp-router)# neighbor 40.1.1.2 remote-as 200
NetIron(config-bgp-router)# network 40.0.0.0 255.0.0.0
NetIron(config-bgp-router)# exit
```

LSP with Redundant Paths

Figure 3.14 shows a signalled LSP configuration that has a primary and a secondary path. The destination for this LSP is 30.1.1.1. The primary path to this destination is through interface POS 2/2, which has a direct link to interface POS 2/2 on R3. If this link should go down, the secondary path will be established. The secondary path goes through R2.

Figure 3.14 LSP configuration with primary and secondary paths

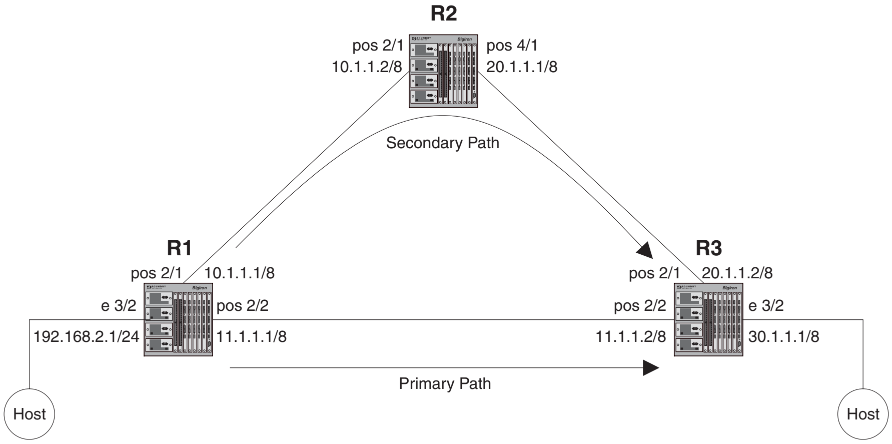

Router R1 is the ingress LER for signalled LSP t3. Packets whose destination is 30.1.1.1 are assigned to this LSP. Two paths are configured, direct_conn and via_r2. Path direct_conn consists of a single strict node, 11.1.1.2, which is a directly connected interface on the destination LSR, R3.

Path via_r2 also consists of a single strict node, 10.1.1.2, a directly connected interface on R2. Since path via_r2 does not specify a node for R3, the hop between R2 and R3 is treated as a hop to a loose node. This means standard hop-by-hop routing is used to determine the path between R2 and R3.

Path direct_conn is the primary path for LSP t3, and path via_r2 is the secondary path. When the LSP is enabled, RSVP signalling messages set up path direct_conn. Packets assigned to this LSP will use this path to reach the destination.

If path direct_conn should fail, path via_r2 is set up, and packets assigned to LSP t3 will then use path via_r2 to reach the destination. By default, the secondary path is not set up until the primary path fails. If you use the **standby** parameter in the configuration of the secondary path, both the primary and secondary paths are set up at the same time, although packets assigned to the LSP travel only down the primary path. If the primary path should fail, the secondary path immediately takes over.

Router R1

The following commands configure Router R1 in Figure 3.14:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# path direct_conn
NetIron(config-mpls-path)# strict 11.1.1.2
NetIron(config-mpls-path)# exit

NetIron(config-mpls)# path via_r2
NetIron(config-mpls-path)# strict 10.1.1.2
NetIron(config-mpls-path)# exit

NetIron(config-mpls)# lsp t3
NetIron(config-mpls-lsp)# to 30.1.1.1
NetIron(config-mpls-lsp)# primary direct
NetIron(config-mpls-lsp)# secondary via_r2
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.1 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit
NetIron(config-mpls)# interface pos 2/2
NetIron(config-posif-2/2)# ip address 11.1.1.1 255.0.0.0
NetIron(config-posif-2/2)# ip ospf area 1
NetIron(config-posif-2/2)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 192.168.2.1 255.255.255.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R2

In the configuration in Figure 3.14, Router R2 serves as a transit LSR for path via_r2. Since path via_r2 is the secondary path for the LSP, it will be used only if the primary path fails.

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface pos 4/1
NetIron(config-posif-4/1)# ip address 20.1.1.1 255.0.0.0
NetIron(config-posif-4/1)# ip ospf area 1
NetIron(config-posif-4/1)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R3

In the configuration in Figure 3.14, Router R3 is the egress LER for LSP t3.

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 20.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface pos 2/2
NetIron(config-posif-2/2)# ip address 11.1.1.2 255.0.0.0
NetIron(config-posif-2/2)# ip ospf area 1
NetIron(config-posif-2/2)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 30.1.1.1 255.0.0.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Static LSP Configuration

Figure 3.15 depicts a static LSP consisting of three LSRs. In a static LSP, label values and inbound/outbound interfaces for labelled packets are specified manually on each LSR.

Figure 3.15 Static LSP configuration

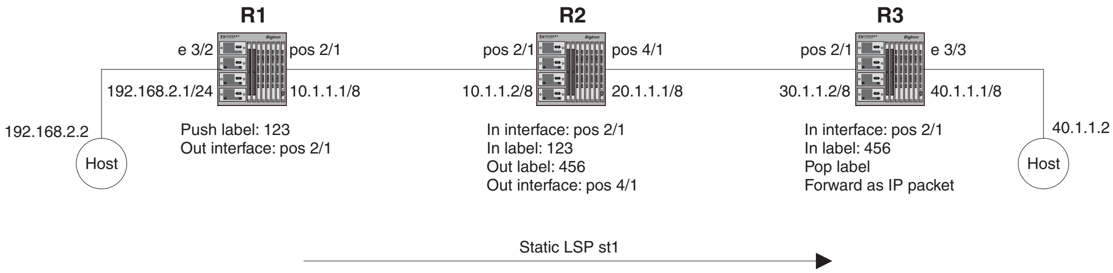

Router R1 is the ingress LER for static LSP st1. When a packet destined for 40.1.1.2 enters the router, it is assigned to the LSP. The Foundry device applies a label with a value of 123 to the packet and forwards the packet out interface POS 2/1 to R2.

R2 is a transit LSR in this LSP. When the LSR receives an MPLS packet on interface POS 2/1, it looks up the label in its internal MPLS forwarding table, which maps the label and inbound interface to a new label and outbound interface. In this configuration, inbound interface POS 2/1 and label 123 is mapped to outbound interface 4/1 and label 456. The LSR replaces label 123 with label 456 and sends the packet out interface POS 4/1 to R3

R3 is the egress LER in this LSP. The egress LER receives labelled packets on interface POS 2/1 and pops the label. The packet can then be forwarded to its destination based on its IP header information.

Router R1

The following commands configure Router R1 in Figure 3.15:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# static-lsp st1
NetIron(config-mpls-static-lsp)# to 40.1.1.2
NetIron(config-mpls-static-lsp)# out-segment pos 2/1 out-label 123
NetIron(config-mpls-static-lsp)# enable
NetIron(config-mpls-static-lsp)# exit
```

NOTE: Prior to issuing the **enable** command on the ingress LER for a static LSP, you should make sure you have already configured and enabled the LSP on the egress LER and on any transit LSRs. When the LSP is enabled on the ingress LER, packets can be assigned to the LSP and forwarded over the configured path of LSRs.

```
NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.1 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/2
NetIron(config-if-e100-3/2)# ip address 192.168.2.1 255.255.255.0
NetIron(config-if-e100-3/2)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R2

The following commands configure Router R2 in Figure 3.15:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# static-lsp st1
NetIron(config-mpls-static-lsp)# in-segment pos 2/1 in-label 123
NetIron(config-mpls-static-lsp)# out-segment pos 4/1 out-label 456
NetIron(config-mpls-static-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 10.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface pos 4/1
NetIron(config-posif-2/1)# ip address 20.1.1.1 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Router R3

The following commands configure Router R3 in Figure 3.15:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# static-lsp st1
NetIron(config-mpls-static-lsp)# in-segment pos 2/1 in-label 456
NetIron(config-mpls-static-lsp)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# ip address 30.1.1.2 255.0.0.0
NetIron(config-posif-2/1)# ip ospf area 1
NetIron(config-posif-2/1)# exit

NetIron(config-mpls)# interface e 3/3
NetIron(config-if-e100-3/3)# ip address 40.1.1.1 255.0.0.0
NetIron(config-if-e100-3/3)# exit

NetIron(config)# router ospf
NetIron(config-ospf-router)# area 1
NetIron(config-ospf-router)# exit
```

Chapter 4
Configuring LDP

Foundry devices support Label Distribution Protocol (LDP) for setting up non-traffic-engineered tunnel LSPs in an MPLS network. LDP is described in RFC 3036.

LDP Overview

When used to create tunnel LSPs, LDP allows a set of destination IP prefixes (known as a Forwarding Equivalence Class or FEC) to be associated with an LSP. Each LSR establishes a peer relationship with its neighboring LDP-enabled routers and exchanges label mapping information. This label mapping information is stored in an LDP database on each LSR. When an LSR determines that one of its peers is the next hop for a FEC, it uses the label mapping information from the peer to set up an LSP that is associated with the FEC. It then sends label mapping information to its upstream peers, allowing the LSP to extend across the MPLS network.

Foundry devices advertise their loopback addresses to their LDP peers as a 32-bit “prefix” type FEC. When an LSR installs a label for a FEC, it also creates an MPLS tunnel route, which is then made available to routing applications. This allows each router to potentially be an ingress LER for an LSP whose destination is the Foundry device's loopback address.

The result of an LDP configuration is a full mesh of LSPs in an MPLS network, with each LDP-enabled router a potential ingress, transit, or egress LSR, depending on the destination.

Foundry’s implementation supports the following aspects of LDP:

Liberal Label Retention – Each LSR sends its peers Label Mapping messages, which map a label to a FEC. Peer LSR receiving these messages retain all of the mappings, even though they may not actually be used for data forwarding.

Unsolicited Label Advertisement – The LSR sends Label Mapping messages to its LDP peers even though they did not explicitly request them.

Ordered Label Distribution – The LSR sends a Label Mapping message to its peers only when it knows the next hop for a FEC, or is itself an egress LER for the FEC. If an LSR does not know the next hop for a FEC, and is not an egress LER for the FEC, it waits until a downstream LSR sends it a Label Mapping message for the FEC. At this point, the LSR can send Label Mapping messages for the FEC to its peers. This allows label mappings to be distributed, in an orderly fashion, starting from the egress LER and progressing upstream.

Configuring LDP on an Interface

To use LDP, a loopback address (with a 32-bit mask) ***must*** be configured on the LSR. The first loopback address configured on the device is used in its LDP identifier. If the loopback address used in the LDP identifier is removed, all LDP functions on the LSR will be shut down. LDP sessions between the LSR and its peers will be terminated, and LDP-created tunnels will be removed. If other loopback interfaces are configured on the device, the lowest-numbered loopback address will then be used as a new LDP identifier. LDP sessions and tunnels will be set up using this new LDP identifier.

To configure LDP on an interface, enter commands such as the following:

```
NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface e 1/2
NetIron(config-mpls)# ldp-enable
```

Syntax: ldp-enable

NOTE: You should enable LDP on the same set of interfaces that IGP routing protocols such as OSPF and IS-IS are enabled.

Setting the LDP Hello Interval

The LDP hello interval controls how often the device sends out LDP Hello messages. Hello messages are used to maintain LDP sessions between the device and its LDP peers. You can set the interval for LDP Link Hello messages (LDP Hello messages multicast to all routers on the sub-net), as well as LDP Targeted Hello messages (LDP Hello messages unicast to a specific address, such as a VLL peer).

The default interval for LDP Link Hello messages is 5 seconds. To set the interval for LDP Link Hello messages to 10 seconds, enter the following command:

```
NetIron(config-mpls)# ldp
NetIron(config-mpls-ldp)# hello-interval 10
```

Syntax: ldp

Syntax: hello-interval <seconds>

The default interval for LDP Targeted Hello messages is 15 seconds. To set the interval for LDP Targeted Hello messages to 20 seconds, enter the following command:

```
NetIron(config-mpls)# ldp
NetIron(config-mpls-ldp)# hello-interval target 20
```

Syntax: hello-interval target <seconds>

The LDP hello interval can be from 1 – 65535 seconds. When you set a new LDP hello interval, it takes effect immediately.

Setting the LDP Hold Time

The LDP hold time specifies how long the device waits for its LDP peers to send a Hello message. If the device does not receive a Hello message within this time, the LDP session with the peer can be terminated. The device includes the hold time in the Hello messages it sends out to its LDP peers. You can set the hold time sent in LDP Link Hello messages and LDP Targeted Hello messages.

The default hold time included in LDP Link Hello messages is 15 seconds. To set the hold time included in LDP Link Hello messages to 20 seconds, enter the following command:

```
NetIron(config-mpls)# ldp
NetIron(config-mpls-ldp)# hello-timeout 20
```

Syntax: hello-timeout <seconds>

The default hold time included in LDP Targeted Hello messages is 45 seconds. To set the hold time included in LDP Targeted Hello messages to 60 seconds, enter the following command:

```
NetIron(config-mpls)# ldp
NetIron(config-mpls-ldp)# hello-timeout target 60
```

Syntax: hello-timeout target <seconds>

The LDP hold time can be from 1 – 65535 seconds. When you set a new LDP hold time, it takes effect for new Targeted peers or when LDP is restarted.

The hold time setting applies only to the hold time the device sends out to peers; it does not affect the hold time the device uses to time out those peers. This is determined from the hold time value sent by the peers to the device.

NOTE: The LDP hold time must always be greater than twice the LDP hello interval.

Displaying LDP information

You can display the following information about LDP:

- The LDP version number, as well as the LSR's LDP identifier and loopback number
- Information about active LDP-created LSPs on the device
- Information about LDP-created tunnel LSPs for which this device is the ingress LER
- The contents of the LDP database
- Information about the LDP session between this LSR and its LDP peers
- Information about the connection between this LSR and its LDP peers
- Information about LDP-enabled Interfaces on the LSR

Displaying the LDP Version

To display the LDP version number, the LSR's LDP identifier and loopback number, and the LDP hello interval and hold time:

```
NetIron(config)# show mpls ldp
Label Distribution Protocol version 1
  LSR ID: 2.2.2.2, using Loopback 1 (deleting it will stop LDP)
  Hello interval: Link 5 sec, Targeted 15 sec
  Hold time value sent in Hellos: Link 15 sec, Targeted 45 sec
```

Syntax: show mpls ldp

Table 4.1 lists the information displayed by the **show mpls ldp** command.

Table 4.1: Output from the show mpls ldp command

This Field...	Displays...
Label Distribution Protocol version	The LDP version.
LSR ID:	The LDP identifier of the device and the loopback interface number being used by LDP. LDP advertises the address of this loopback interface in Address messages.
Hello interval:	How often the device sends out LDP Link Hello and Targeted Hello messages.
Hold time value sent in Hellos:	How long the device waits for its LDP peers to send a Hello message. The hold time is included in the Link Hello and Targeted Hello messages it sends out to its LDP peers.

Displaying Information about LDP-Created LSPs

You can display information about active LDP-created LSPs for which this device is an ingress, transit or egress LSR. For example:

```
NetIron(config)# show mpls ldp path
Upstr-session(label)      Downstr-session(label, intf)    Destination route
33.3.3.3:0(3)             (egress)                        11.1.1.1/32
22.2.2.2:0(3)             (egress)                        11.1.1.1/32
33.3.3.3:0(1024)          22.2.2.2:0(3, e2/10)            22.2.2.2/32
22.2.2.2:0(1024)          22.2.2.2:0(3, e2/10)            22.2.2.2/32
(ingress)                 22.2.2.2:0(3, e2/10)            22.2.2.2/32
33.3.3.3:0(1026)          33.3.3.3:0(3, e2/20)            33.3.3.3/32
22.2.2.2:0(1026)          33.3.3.3:0(3, e2/20)            33.3.3.3/32
(ingress)                 33.3.3.3:0(3, e2/20)            33.3.3.3/32
```

Syntax: show mpls ldp path

Each line in the output of the **show mpls ldp path** command shows information about an LSP created through LDP. The command lists the incoming and outgoing labels applied to packets in each LSP. For example, the third line in the output above indicates that MPLS packets received from upstream peer 33.3.3.3 with label 1025 are to be transmitted to downstream peer 22.2.2.2 with label 3.

Note that in this context, “upstream” and “downstream” refer to the direction that data traffic flows in an LSP. This is opposite of the direction that labels are distributed using LDP.

Additionally, the output of this command indicates that the device has received a label for the destination IP prefix (that is, the attached route) from the downstream peer and then advertised a label for that IP prefix to the upstream peer. For example, the third line shows that the device has received a Label Mapping message binding label 3 to IP prefix 22.2.2.2/32 from its downstream peer 22.2.2.2, and the device has sent a Label Mapping message binding label 1025 to IP prefix 22.2.2.2/32 to upstream peer 33.3.3.3.

Table 4.2 lists the information displayed by the **show mpls ldp path** command.

Table 4.2: Output from the show mpls ldp path command

This Field...	Displays...
Upstr-session(label)	The LDP identifier of the upstream peer, as well as the incoming label. Note that upstream session information does not apply to LSPs for which this is the ingress LER. Since the device uses a per-platform label space, the incoming interface for LDP-created is not relevant.
Downstr-session(label, intf)	The LDP identifier of the downstream peer, as well as the outgoing label and interface. Note that downstream session information does not apply to LSPs for which this is the egress LER.
Attached route	The destination route bound to this LSP.

Displaying LDP Tunnel LSP Information

To display information about LDP-created LSPs for which this device is the ingress LER, enter the following command:

```
NetIron# show mpls ldp tunnel
                Oper      Tunnel    Outbound
To              State     Intf      Intf
22.2.2.2        UP        tnl0      p3/1
```

```
33.3.3.3        UP        tnl1      p3/2
```

Syntax: show mpls ldp tunnel

The following table describes the output of the **show mpls ldp tunnel** command.

Table 4.3: Output from the show mpls ldp tunnel command

This Field...	Displays...
To	The egress LER for the LSP.
Oper State	The operational state of the LSP. This field indicates whether the LSP has been established through LDP signalling and is capable of having packets forwarded through it.
Tunnel Intf	The MPLS tunnel interface port ID.
Outbound Intf	The outbound interface for the LSP.

Displaying the Contents of the LDP Database

You can display the contents of the LSR's LDP Label Information Base. This database contains all the labels it has learned from each of its LSR peers, as well as all of the labels it has sent to its LDP peers.

```
NetIron# show mpls ldp database
Session 13.13.13.13:0 - 11.11.11.11:0
 Downstream label database:
   Label     Prefix              State
   3         11.11.11.11/32      Installed
   1066      13.13.13.13/32      Retained
   1227      12.12.12.12/32      Retained
   1228      14.14.14.14/32      Retained
 Upstream label database:
   Label     Prefix
   3         13.13.13.13/32
   1024      14.14.14.14/32
   1027      12.12.12.12/32
   1028      11.11.11.11/32
```

Syntax: show mpls ldp database

For each LDP session, the **show mpls ldp database** command displays the following information:

Table 4.4: Output from the show mpls ldp database command

This Field...	Displays...
Session	The LDP identifiers of this LSR and its peer.
Downstream label database:	Information about labels received from the LDP peer
Upstream label database:	Information about labels distributed by this LSR to the LDP peer. The device sends the same label for a given prefix to all of its upstream peers. In the example above, label 3 is mapped to prefix 14.14.14.14/32. This device will send this label mapping to each of its upstream peers.
Label	The label value received from or distributed to LDP peers
Prefix	The destination route associated with the label.

Table 4.4: Output from the show mpls ldp database command (Continued)

This Field...	Displays...
State	Whether the label is actively being used for data forwarding. This can be one of the following: “Installed” indicates that the label is being used with an active LDP-created LSP to forward packets. “Retained” indicates that the label is not being used for packet forwarding. Since Foundry LSRs use Liberal Label Retention, these unused labels are retained in the database and not discarded.

Displaying LDP Session Information

To display information about the LDP session between this LSR and its LDP peers, enter the following command:

```
NetIron# show mpls ldp session
Peer LDP ID         State               Role      Hold time
22.2.2.2:0          Operational         passive   34

Peer LDP ID              State          Targeted  My Role   Max Hold  Time Left
1.1.1.1:0                Operational    No        Active    36        26
5.5.5.5:0                Operational    No        Passive   36        26
4.4.4.4:0                Operational    Yes       Passive   36        26
```

Syntax: show mpls ldp session [detail]

The following table describes the output of the **show mpls ldp session** command.

Table 4.5: Output from the show mpls ldp session command

This Field...	Displays...
Peer LDP ID	The LDP identifier of the peer LSR. The first four octets identify the peer LSR; the second two octets identify a label space on the LSR. For LSRs that use per-platform label spaces, the second two octets are always zero.
State	The current state of the LDP session between this LSR and its peer, as defined in RFC 3036. This can be "Nonexistent", "Initialized", "OpenRec", "OpenSent", or "Operational".
Targeted	Whether the session was established using Targeted Hello messages (that is, through extended discovery).
My Role	Whether this LSR is playing the “active” or “passive” role in LDP session establishment (as defined in RFC 3036). The LSR with the higher LSR ID plays the active role in LDP session establishment.
Max Hold	The number of seconds that the “Hold time remain” counter is reset to once a KeepAlive message is received from the peer.
Time Left	The amount of time, in seconds, before the LDP session times out if no KeepAlive message is received from the peer.

To display more detailed information about the LDP session between this LSR and its LDP peers, enter the following command:

```
NetIron# show mpls ldp session
Peer LDP ID: 1.1.1.1:0, Local LDP ID: 2.2.2.2:0, State: Operational
  Targeted: No, Role: Active, Next keepalive: 2 sec, Hold time left: 26 sec
  Keepalive interval: 6 sec, Max hold time: 36 sec
```

```
    Neighboring interfaces: p4/1
    TCP connection: 2.2.2.2:9002--1.1.1.1:646, State: ESTABLISHED
    Next-hop addresses received from the peer:
      1.1.1.1  10.1.1.1  11.1.1.1  12.1.1.1  13.1.1.1  40.1.1.1  43.1.1.1
```

Syntax: show mpls ldp session detail

For each established LDP session, the command displays the following information:

Table 4.6: Output from the show mpls ldp session detail command

This Field...	Displays...
Peer LDP ID:	The LDP identifier of the peer LSR. The first four octets identify the peer LSR; the second two octets identify a label space on the LSR. For LSRs that use per-platform label spaces, the second two octets are always zero.
Local LDP ID:	This LSR’s LDP identifier.
State:	The LDP session state, as defined in RFC 3036. This can be "Nonexistent", "Initialized", "OpenRec", "OpenSent", or "Operational".
Targeted:	Whether the session was established using Targeted Hello messages (that is, through extended discovery).
Role:	Whether this LSR is playing an “active” or “passive” role in session establishment.
Next keepalive:	If this LDP session is established, the amount of time, in seconds, before the next KeepAlive message is sent to the active peer. If this LSR is the active peer, prior to establishing a session with the passive peer, the text "Next Initialization:" is displayed instead. The "Next Initialization:" value indicates, in seconds, when the next Initialization message will be sent to the passive peer.
Hold time left:	The amount of time, in seconds, before the LDP session times out if no KeepAlive message is received from the peer.
Max hold time:	The number of seconds that the “Hold time remain” counter is reset to once a KeepAlive message is received from the peer.
Keepalive interval:	The amount of time the LSR waits for an LDP PDU from the peer. If this amount of time passes without receiving an LDP PDU from the peer, the LDP session is terminated.
Neighboring interfaces:	The interfaces where an LDP neighbor/adjacency relationship has been established with the peer. If there are multiple connections between two LDP-enabled peers, there can be multiple neighboring interfaces.
TCP connection:	The local and remote IP addresses and port numbers for the TCP connection between the peers.
State:	The state of the TCP connection between the peers.
Next-hop addresses received from the peer:	The next-hop addresses received from the peer in LDP address messages. The LSR uses this list of addresses to determine whether the peer is the correct next hop for a destination route. If one of the addresses in this list is the correct next hop for the route, the label received from the peer is installed for that route, allowing it to be used for data forwarding.

Displaying LDP Neighbor Connection Information

To display information about the connection between this LSR and its LDP-enabled neighbors, enter the following command:

```
NetIron# show mpls ldp neighbor
Nbr Transport        Interface      Nbr LDP ID          Max Hold  Time Left
1.1.1.1              p4/1           1.1.1.1:0           15        14
5.5.5.5              p3/2           5.5.5.5:0           15        11
4.4.4.4              (targeted)     4.4.4.4:0           15        13
```

Syntax: show mpls ldp neighbor

Table 4.7: Output from the show mpls ldp neighbor command

This Field...	Displays...
Nbr Transport	The transport address of the LDP neighbor.
Interface	The interface to which the LDP neighbor is connected. “(targeted)” indicates that the session between this device and the neighbor was established using Targeted Hello messages (that is, through extended discovery).
Nbr LDP ID	The neighbor’s LDP identifier
Max Hold	The number of seconds the device waits for its LDP peers to send a Hello message.
Time Left	The amount of time, in seconds, before the LDP neighbor times out if no Hello message is received from the neighbor.

Displaying Information about LDP-Enabled Interfaces

To display information about the LDP enabled interfaces on the LSR, enter the following command:

```
NetIron# show mpls ldp interface
                     Label-space    Nbr            Hello          Next
Interface            ID             Count          Interval       Hello
p4/1                 0              1              5              3
(targeted)           1              0              15             3
```

Syntax: show mpls ldp interface

Table 4.8: Output from the show mpls ldp interface command

This Field...	Displays...
Interface	The slot and port number of the LDP-connected interface. “(targeted)” shows information about unicast Targeted Hello messages sent to VLL peers.
Label-space ID	The label space ID. For LSRs that use per-platform label spaces, the second two octets are always zero.
Nbr Count	The number of LDP peers/adjacencies that has been established on this interface. This number can be greater than 1 if this is a multi-access network.

Table 4.8: Output from the show mpls ldp interface command (Continued)

This Field...	Displays...
Hello Interval	The number of seconds between LDP Hello messages.
Next Hello	The number of seconds before the next LDP Hello message is sent (multicast) to the interface. For a targeted interface, the LDP Hello message is unicast.

Sample LDP Configurations

Figure 4.1 illustrates a sample configuration with three LDP-enabled LSRs.

Figure 4.1 Sample LDP configuration

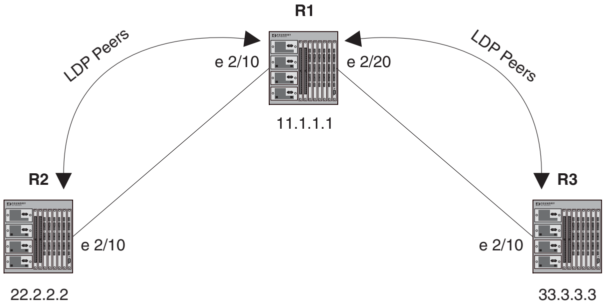

Router R1

The following commands configure Router R1 in Figure 4.1:

```
R1(config)# interface loopback 1
R1(config-lbif-1)# ip address 11.1.1.1/32
R1(config-lbif-1)# exit
```

```
R1(config)# router mpls
R1(config-mpls)# mpls-interface e 2/10
R1(config-mpls)# ldp-enable
R1(config-mpls)# mpls-interface e 2/20
R1(config-mpls)# ldp-enable
R1(config-mpls)# exit
```

```
R1(config)# ip route 22.2.2.2/32 10.1.1.2
R1(config)# ip route 33.3.3.3/32 20.1.1.2
R1(config)# route-only
```

```
R1(config)# interface ethernet 2/10
R1(config-if-2/10)# enable
R1(config-if-2/10)# ip address 10.1.1.1/24
R1(config-if-2/10)# exit
```

```
R1(config)# interface ethernet 2/20
R1(config-if-2/20)# enable
R1(config-if-2/20)# ip address 20.1.1.1/24
```

Router R2

The following commands configure Router R2 in Figure 4.1:

```
R2(config)# interface loopback 1
R2(config-lbif-1)# ip address 22.2.2.2/32
R2(config-lbif-1)# exit
```

```
R2(config)# router mpls
R2(config-mpls)# mpls-interface e 2/10
R2(config-mpls)# ldp-enable
R2(config-mpls)# exit

R2(config)# ip route 11.1.1.1/32 10.1.1.1
R2(config)# ip route 33.3.3.3/32 10.1.1.1
R2(config)# route-only

R2(config)# interface ethernet 2/20
R2(config-if-2/20)# enable
R2(config-if-2/20)# ip address 10.1.1.2/24
R2(config-if-2/20)# exit
```

Router R3

The following commands configure Router R3 in Figure 4.1:

```
R3(config)# interface loopback 1
R3(config-lbif-1)# ip address 33.3.3.3/32
R3(config-lbif-1)# exit

R3(config)# router mpls
R3(config-mpls)# mpls-interface e 2/10
R3(config-mpls)# ldp-enable
R3(config-mpls)# exit

R3(config)# ip route 11.1.1.1/32 20.1.1.1
R3(config)# ip route 22.2.2.2/32 20.1.1.1
R3(config)# route-only

R3(config)# interface ethernet 2/20
R3(config-if-2/20)# enable
R3(config-if-2/20)# 20.1.1.2/24
R3(config-if-2/20)# exit
```

Sample LDP Configuration with VLL

Figure 4.2 illustrates a sample Virtual Leased Line (VLL) configuration that uses LDP tunnel LSPs.

Figure 4.2 MPLS VLL configuration with LDP tunnel LSPs

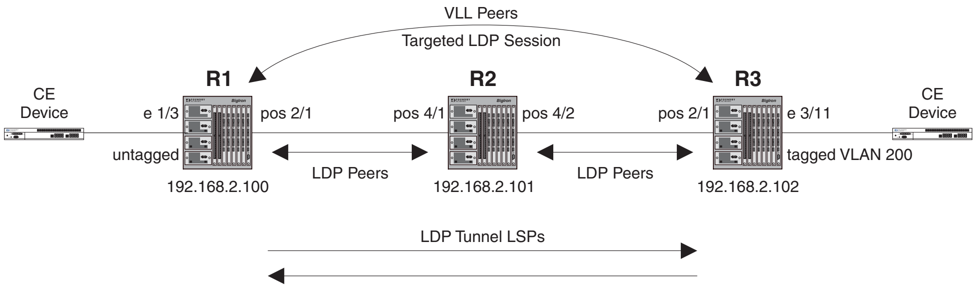

In this example, routers R1 and R3 are Provider Edge (PE) routers configured as VLL peers. R1 and R3 have established a targeted LDP session to exchange VLL label information. When this targeted LDP session is established, each router advertises its locally assigned VC label and VC ID to its VLL peer.

In addition, LDP sessions have been established between R1 – R2 and R2 – R3. LDP tunnel LSPs exist in each direction between R1 and R3. When the CE device forwards a Layer 2 packet to R1, the router assigns the packet to an LSP whose destination is R3. R1 encapsulates the packet as an MPLS packet, adding a tunnel label and the VC label advertised to the router by R3. The MPLS packet is then forwarded over the outbound interface indicated by the tunnel label to the next hop in the LSP.

When the MPLS packet reaches R2, the penultimate LSR in the tunnel LSP, R2 pops the tunnel label, leaving the packet with only the VC label, then forwards the packet to R3.

R3 examines the VC label in the packet. On R3, the VC label is mapped to the user-specified endpoint for the VLL. In this example, the endpoint consists of VLAN ID 200 and interface 3/11. R3 then pops the VC label, tags the Layer 2 packet with VLAN 200, then forwards the packet out interface 3/11.

In the opposite direction, R3 assigns traffic received from the CE device to a tunnel LSP destined for R1, pushes tunnel and VC labels onto the packets, and forwards them to the next hop in the LSP. When the packets reach R1, the router pops the VC label and forwards the Layer 2 packets out the interface indicated by the VLL endpoint. In this example, the endpoint consists of interface 1/3, so the packets are forwarded untagged out interface 1/3 to the CE device.

Router R1

The following commands configure Router R1 in Figure 4.2:

```
R1(config-mpls)# interface loopback 1
R1(config-lbif-1)# port-name Generic All-Purpose Loopback
R1(config-lbif-1)# ip address 192.168.2.100/32
R1(config-lbif-1)# ip ospf area 0
R1(config-lbif-1)# exit

R1(config)# router mpls
R1(config-mpls)# mpls-interface all-pos
R1(config-mpls)# ldp-enable
R1(config-mpls)# exit

R1(config-mpls)# vll VLL_to_R3 40000
R1(config-mpls-vll)# vll-peer 192.168.2.102
R1(config-mpls-vll)# untagged e 1/3
R1(config-mpls-vll)# exit

R1(config)# ip router-id 192.168.2.100
R1(config)# router ospf
R1(config-ospf-router)# area 0
R1(config-ospf-router)# exit

R1(config-mpls)# interface e 1/3
R1(config-if-e100-1/3)# port-name VLL_endpoint
R1(config-if-e100-1/3)# enable
R1(config-if-e100-1/3)# exit

R1(config-mpls)# interface pos 2/1
R1(config-posif-2/1)# port-name Connection_to_R2
R1(config-posif-2/1)# enable
R1(config-posif-2/1)# ip address 192.168.37.1/30
R1(config-posif-2/1)# ip ospf area 0
R1(config-posif-2/1)# exit
```

Router R2

The following commands configure Router R2 in Figure 4.2:

```
R2(config-mpls)# interface loopback 1
R2(config-lbif-1)# port-name Generic All-Purpose Loopback
R2(config-lbif-1)# ip address 192.168.2.101/32
R2(config-lbif-1)# ip ospf area 0
R2(config-lbif-1)# exit

R2(config)# router mpls
R2(config-mpls)# mpls-interface all-pos
R2(config-mpls)# ldp-enable
R2(config-mpls)# exit
```

```
R2(config)# ip router-id 192.168.2.101
R2(config)# router ospf
R2(config-ospf-router)# area 0
R2(config-ospf-router)# exit

R2(config-mpls)# interface pos 4/1
R2(config-posif-4/1)# enable
R2(config-posif-4/1)# ip address 192.168.40.1/30
R2(config-posif-4/1)# ip ospf area 0
R2(config-posif-4/1)# exit

R2(config-mpls)# interface pos 4/2
R2(config-posif-4/2)# enable
R2(config-posif-4/2)# ip address 192.168.40.9/30
R2(config-posif-4/2)# ip ospf area 0
R2(config-posif-4/2)# exit
```

Router R3

The following commands configure Router R3 in Figure 4.2:

```
R3(config-mpls)# interface loopback 1
R3(config-lbif-1)# port-name Generic All-Purpose Loopback
R3(config-lbif-1)# ip address 192.168.2.102/32
R3(config-lbif-1)# ip ospf area 0
R3(config-lbif-1)# exit

R3(config)# router mpls
R3(config-mpls)# mpls-interface all-pos
R3(config-mpls)# ldp-enable
R3(config-mpls)# exit

R3(config-mpls)# vll VLL_to_R1 40000
R3(config-mpls-vll)# vll-peer 192.168.2.100
R3(config-mpls-vll)# vlan 200
R3(config-mpls-vll-vlan)# tagged e 3/11
R3(config-mpls-vll-vlan)# exit
R3(config-mpls-vll)# exit

R3(config)# ip router-id 192.168.2.102
R3(config)# router ospf
R3(config-ospf-router)# area 0
R3(config-ospf-router)# exit

R3(config-mpls)# interface e 3/11
R3(config-if-e100-3/11)# port-name VLL_endpoint
R3(config-if-e100-3/11)# enable
R3(config-if-e100-3/11)# exit

R3(config-mpls)# interface pos 2/1
R3(config-posif-2/1)# port-name Connection_to_R2
R3(config-posif-2/1)# enable
R3(config-posif-2/1)# ip address 192.168.41.1/30
R3(config-posif-2/1)# ip ospf area 0
R3(config-posif-2/1)# exit
```

Chapter 5
Configuring MPLS Virtual Leased Line

This chapter explains how to configure MPLS ***Virtual Leased Line (VLL)*** on a Foundry device. MPLS VLL is a method for providing point-to-point Ethernet/VLAN connectivity over an MPLS domain. This functionality is outlined in the IETF documents “draft-martini-l2circuit-trans-mpls-07.txt” and “draft-martini-l2circuit-encap-mpls-03.txt”.

Overview

This chapter is divided into the following sections:

- “How MPLS VLL Works” on page 5-1 describes how packets are encapsulated and forwarded over an MPLS VLL.
- “Configuring MPLS VLLs” on page 5-5 describes how to set up MPLS VLLs on Foundry devices using the Command Line Interface (CLI).
- “Displaying MPLS VLL Information” on page 5-8 describes the commands used to display information about an MPLS VLL configuration.
- “Sample MPLS VLL Configuration” on page 5-11 illustrates a sample MPLS VLL configuration and lists the CLI commands used for implementing it.

How MPLS VLL Works

The following diagram illustrates how packets are forwarded over an MPLS VLL.

Figure 5.1 Forwarding packets over an MPLS VLL

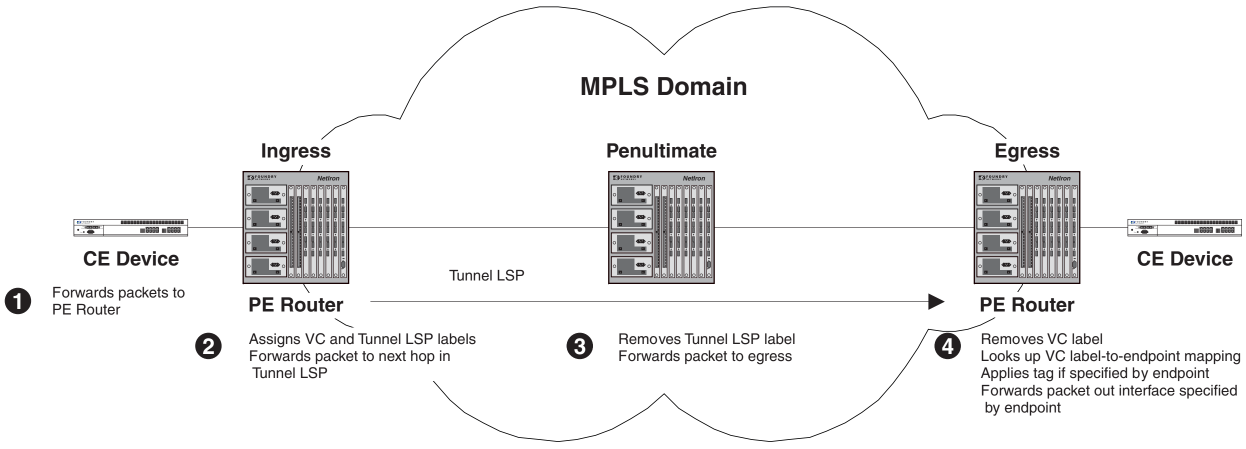
MPLS Domain
Ingress
Penultimate
Egress
CE Device
CE Device
Tunnel LSP
PE Router
PE Router
1 Forwards packets to PE Router
2 Assigns VC and Tunnel LSP labels
Forwards packet to next hop in Tunnel LSP
3 Removes Tunnel LSP label
Forwards packet to egress
4 Removes VC label
Looks up VC label-to-endpoint mapping
Applies tag if specified by endpoint
Forwards packet out interface specified by endpoint

Packets are forwarded over an MPLS VLL as follows:

1. A Customer Edge (CE) device forwards a packet to an LER serving as a Provider Edge (PE) router at the edge of the MPLS domain.

2. The PE router assigns the packet to an RSVP-signalled LSP whose destination is an LER (also serving as a PE router) that is connected to a CE device at the other end of the MPLS domain. The PE router at the other end of the MPLS domain is known as this PE router's ***VLL peer***. The RSVP-signalled LSP used to reach the VLL peer is known as the ***tunnel LSP***.

   If a Class of Service (COS) value is set for the VLL, the Foundry device selects a tunnel LSP that also has this COS value, if one is available. If no tunnel LSP with this COS value is available, the Foundry device selects a tunnel LSP with the highest configured COS value (although never higher than the COS setting for the VLL). See "QoS for VLL Traffic" on page 5-3 for more information.

   If there are multiple tunnel LSPs that can be used to reach the VLL peer, the PE router selects one of them using a round-robin method.

   The PE router pushes two labels onto the packet:

   - The inner ***VC label*** is used for determining what happens to the packet once it reaches the VLL peer. This label is significant only to the VLL peer.
   - The outer ***tunnel label*** is used for forwarding the packet through the MPLS domain. This label corresponds to an RSVP-signalled tunnel LSP.

   See "MPLS VLL Packet Encoding" on page 5-3 for information on the structure of packets forwarded along an MPLS VLL. After applying the two labels to the packet, the PE router forwards it to the next LSR in the tunnel LSP.

3. The penultimate LSR in the tunnel LSP removes the tunnel label and forwards the packet (now with the VC label as the top label) to the PE router at the other edge of the MPLS domain.

4. The VLL peer at the egress of the tunnel LSP examines the VC label. On the VLL peer, the VC label is mapped to an ***endpoint*** for the VLL. The endpoint of a VLL specifies what happens to packets exiting the VLL.

   The endpoint can specify either an untagged port or a tagged port. For untagged ports, the endpoint consists of an interface. For tagged ports the endpoint consists of a VLAN ID and an interface. The egress LER removes the VC label, and forwards the packet out the interface specified as the endpoint. If the endpoint is a tagged port, the device transmits the packet with the specified VLAN ID, forwarding it out the specified interface to the CE device.

The two VLL peers advertise VC labels to each other using the ***Label Distribution Protocol (LDP)***. Each PE router attempts to initiate an LDP session with its VLL peer. Once the LDP session is established, the locally assigned VC label, along with a VLL VC ID, is advertised to the VLL peer. In a similar way, the PE also learns the remotely assigned VC label from the VLL peer. Alternatively, you can configure static local and remote VC labels manually on both VLL peers; in this case, LDP is not used.

MPLS VLLs are not involved with spanning tree operations.

MPLS VLL Packet Encoding

When a packet is forwarded from the CE device, the PE router encapsulates it as an MPLS packet, applying two labels. The resulting MPLS packet has the following structure:

Figure 5.2 Structure of a packet forwarded over an MPLS VLL

MPLS-PPP	Tunnel Label	VC Label	Ethernet Header	Ethernet Payload

The MPLS-PPP header and tunnel label are the same as for labeled IP packets. The S bit in the tunnel label is zero, indicating that it is not the bottom of the stack. The VC label is significant only to the PE router at the other end of the VLL.

The Ethernet header may be tagged or untagged. The Ethernet payload may be greater than 1500 bytes if the Super Aggregated VLAN feature is used.

QoS for VLL Traffic

By default, packets travelling through an MPLS domain are treated equally from a QoS standpoint, in a best effort manner. However, if a Layer 2 packet has an internal priority in its 802.1q tag, or the LSP or VLL to which the packet is assigned has a configured Class of Service (COS) value, QoS can be applied to the packet in the MPLS domain. The internal priority or COS value is mapped to a value in the EXP field of the packet’s MPLS header. The value in the EXP field is then mapped to an internal forwarding priority, and the packet is sent to the hardware forwarding queue that corresponds to the internal forwarding priority.

QoS for VLL Traffic at the Ingress LER

The following methods can be used to provide QoS to packets entering a VLL:

- Use the COS value assigned to the tunnel LSP used to reach the VLL peer.

  When a tunnel LSP has a user-configured COS value, all packets in all VLLs travelling through the tunnel LSP receive the same QoS.

- Use the COS value assigned to the VLL.

  If a COS value is set for the VLL, the Foundry device selects a tunnel LSP that also has this COS value, if one is available. If no tunnel LSP with this COS value is available, the Foundry device selects a tunnel LSP with the highest configured COS value (although never higher than the COS setting for the VLL).

  If the selected tunnel LSP does not have a COS value, the VLL’s configured COS value is used to provide QoS. The VLL’s COS value is mapped to a value in the EXP field. This allows traffic multiple VLLs using a single tunnel LSP, traffic from each VLL can receive different QoS treatment.

- Use the priority in the packet’s 802.1q tag.

  When neither the tunnel LSP nor the VLL has a configured COS value, the device examines the priority in the Layer 2 packet's 802.1q tag, if the packet has one. Consequently, Layer 2 packets with the same 802.1q priority receive the same QoS in the VLL.

- Use the configured priority of the port.

  If neither the tunnel LSP nor the VLL has a configured COS value, and the Layer 2 packet does not have an 802.1q priority, QoS can be provided based on the priority of the incoming port. A port can be assigned a priority from 0 (lowest priority) to 7 (highest priority). The default port priority is 0.

  By assigning different priorities to the ports where customer edge (CE) devices are connected (that is, the VLL endpoints), you can provide QoS to untagged Layer 2 traffic received from different customer locations.

When a packet enters a VLL, the PE router, serving as both the VLL endpoint and the ingress of a tunnel LSP pushes two labels onto the packet: the inner VC label and the outer tunnel label. The packet's priority is carried in the EXP field of the MPLS label header. Both the VC label and the tunnel label carry the same value in the EXP field.

The following table lists how a Layer 2 packet's priority is mapped to a value in the EXP field and how the EXP value is mapped to a priority queue.

Tunnel LSP configured COS / VLL configured COS / 802.1q priority / Configured Port Priority	Value placed in the tunnel and VC label EXP field	Priority Queue
6, 7	6	qosp3 (highest priority)
4, 5	4	qosp2
2, 3	2	qosp1
0, 1	0	qosp0 (best effort)

QoS for VLL Traffic at Transit LSRs

At each transit LSR, the device reads the value in the tunnel label's EXP field and places the incoming EXP value in the EXP field of the outbound packet. The outbound MPLS packet is assigned to one of the four priority queues based on the value in the EXP field. The EXP bits in the MPLS header are used to assign the packet to a priority queue as follows:

EXP Bits in Tunnel Label	Priority Queue
6, 7	qosp3 (highest priority)
4, 5	qosp2
2, 3	qosp1
0, 1	qosp0 (best effort)

QoS for VLL Traffic at the Penultimate LSR

When the packet reaches the penultimate LSR in the LSP, its tunnel label is popped, leaving the VC label. The MPLS packet is placed in one of the priority queues using the value in the EXP field of the VC label. Since the VC label has the same EXP value as the tunnel label, the packet is placed in the same queue used for the tunnel LSP.

QoS for VLL Traffic at the Egress LER

At the VLL endpoint, the VC label is popped and the packet is forwarded as a Layer 2 packet. The packet is placed in one of the priority queues based on the contents of the EXP field in the VC label, as follows:

EXP Bits in VC Label	Priority Queue
6, 7	qosp3 (highest priority)
4, 5	qosp2
2, 3	qosp1
0, 1	qosp0 (best effort)

Configuring MPLS VLLs

This section explains how to set up MPLS VLLs on Foundry devices. It contains the following topics:

- “Creating a VLL” on page 5-5
- “Specifying a VLL Peer” on page 5-5
- “Specifying a VLL Endpoint” on page 5-6

Creating a VLL

You create a VLL by entering VLL configuration statements on two PE routers. The two endpoints of a VLL are associated by having the same VLL VC ID on each PE router.

To create an MPLS VLL, enter commands such as the following:

```
NetIron(config-mpls)# vll foundry-sj-to-sf 40000
NetIron(config-mpls-vll)#
```

On the VLL peer (if it is a Foundry device), you would enter commands such as the following:

```
NetIron(config-mpls)# vll foundry-sf-to-sj 40000
NetIron(config-mpls-vll)#
```

Syntax: vll <vll-name> <vll-vc-id> [cos <cos value>]

The <vll-vc-id> corresponds to the user-configurable ID defined in draft-martini-l2circuit-trans-mpls-07.txt.

You can optionally specify a Class of Service (COS) setting for the VLL. If a COS value is set, the Foundry device selects a tunnel LSP that also has this COS value, if one is available. If no tunnel LSP with this COS value is available, the Foundry device selects a tunnel LSP with the highest configured COS value (although never higher than the COS setting for the VLL). The COS value can be between 0 – 7.

Specifying a VLL Peer

The VLL peer is the PE router at the other end of the VLL. As part of VLL configuration, you specify the IP address of the VLL peer.

Each PE router must have tunnel LSP reachability to its VLL peer. Tunnel LSP reachability is defined as having at least one operational LSP tunnel with the destination (the LSP’s “to” address) matching the VLL peer’s IP address. An LSP terminating on the VLL peer but configured with a different destination address would not be considered a match.

If a PE router does not have tunnel LSP reachability to its VLL peer, or if the remote VC label is not yet available, packets from the local interface are discarded at the ingress PE router. If the local interface is administratively disabled or goes down, a VC label withdraw message is sent to the VLL peer.

By default, each PE router attempts to initiate an LDP session through extended discovery with its VLL peer, if a session is not already established. The PE router also allocates a VC label from a per-platform label range that is mapped to the local endpoint. Once the LDP session is established, the locally assigned VC label, along with the VLL VC ID is advertised to the VLL peer in a downstream-unsolicited manner. In a similar way, the PE also learns the remotely assigned VC label from the VLL peer.

Alternatively, you can configure static local and remote VC labels. In this case, no LDP session is established between the VLL peers. Note that if you use static VC labels, you must configure them on both VLL peers manually.

You specify the peer at the other end of the VLL by entering a command such as the following:

```
NetIron(config-mpls-vll)# vll-peer 192.168.2.100
```

Syntax: vll-peer <ip-addr> [<static-local-vc-label> <static-remote-vc-label>]

The IP address of the peer must match that of a destination for a tunnel LSP configured on the device.

Static local and remote VC label values are optional. If configured, <static-local-vc-label> is the VC label value expected for packets forwarded to the local physical port from the VLL peer, and <static-remote-vc-label> is the VC label applied to packets sent to the remote VLL peer.

Acceptable values for <static-local-vc-label> are 800000 – 1048575. If the label value you specify has already been assigned, a message is displayed requesting a different value.

Specifying a VLL Endpoint

The endpoint of a VLL specifies what happens to packets exiting the VLL. You set the endpoint on the local PE router and this endpoint is mapped to a VC label. The VC label is advertised to the remote PE router at the other end of the VLL through LDP. The remote PE router applies this label to packets entering the VLL. When the packet reaches the end of the VLL, the local PE router checks the mapping between the VC label and the endpoint, removes the VC label from the packet, and forwards the packet out the port specified as the endpoint, applying a VLAN tag to the packet if the endpoint specifies a tagged port. A VLL can be tagged on one end and untagged on the other end.

The Customer Edge (CE) device is connected to the PE router over either an untagged port or a tagged port. In the case of a tagged port, each (port, VLAN id) pair is identified as a unique endpoint, and the packets are sent in tagged Ethernet format within the MPLS payload. In the case of an untagged port, an endpoint is identified by the physical port alone, and the packets are sent in untagged Ethernet format within the MPLS payload.

Specifying an Untagged Endpoint

Untagged ports are not associated with any VLAN. A port must be a member of the default VLAN before it can be used in a VLL configuration as an untagged port. Upon configuration as the endpoint of a VLL, the port is taken out of the default VLAN. This means no local broadcast traffic includes this port. A VLL untagged port does not belong to any VLAN. If the port is currently a member of a regular VLAN or another VLL, the configuration attempt should be rejected.

To specify an untagged endpoint for a VLL:

```
NetIron(config-mpls-vll)# untagged e 2/1
```

Syntax: untagged [pos | ethernet] <portnum>

Specifying a Tagged Endpoint

Tagged ports are configured under a VLAN ID. This VLAN ID is only meaningful for this tagged port. Another tagged port may use the same VLAN ID but the two ports are not under the same VLAN.

For tagged ports, a <vlan-id, port> pair constitutes a VLL endpoint. Two tagged ports can have the same VLAN ID, so that service providers do not need to impose specific VLAN IDs on customers. However, there is no local switching among ports with the same VLAN ID, since they are unrelated endpoints.

In this release, you cannot configure a tagged port to belong to both regular VLANs and VLLs. A tagged VLL port can be a part of multiple VLLs, but cannot at the same time belong to a regular VLAN.

As with regular VLANs, if a port is currently a member of a non-default VLAN as an untagged port, it must be returned to the default VLAN before it can be assigned to a VLL as a tagged port.

To specify a tagged endpoint for a VLL:

```
NetIron(config-mpls-vll)# vlan 200
NetIron(config-mpls-vll-vlan)# tagged e 3/11
```

Syntax: vlan <num>

Syntax: tagged [pos | ethernet] <slot/port>

NOTE: In releases prior to 07.6.01, you could not configure a tagged port to belong to a regular VLAN and, at the same time, be the endpoint of a VLL. A tagged VLL port could be a part of multiple VLLs, but could not at the same time belong to a regular VLAN. Release 07.6.01 removed this restriction. A tagged port can now be part of one or more VLLs and at the same time be part of one or more VLANs.

Specifying a Trunk Group as the Endpoint of a VLL

The endpoint of a VLL can be a trunk group. When the endpoint of a VLL is a trunk group, the VLL traffic load is distributed to the customer edge (CE) device across all of the trunk group's ports, using a hashing mechanism based on source and destination MAC addresses.

Figure 5.3 illustrates a sample configuration where a trunk group of two ports serves as the endpoint of a VLL.

Figure 5.3 Specifying a trunk group as the endpoint of a VLL

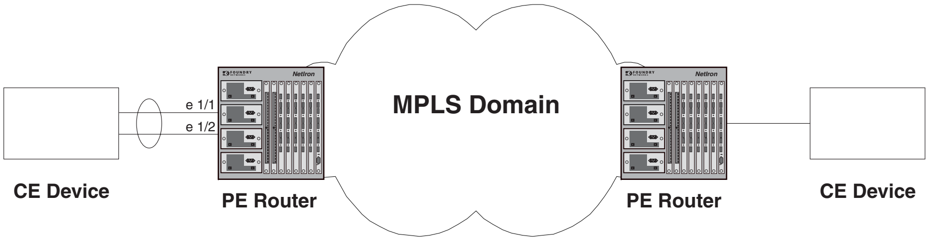

To configure a trunk group like the one in Figure 5.3, enter commands such as the following:

```
NetIron(config)# trunk server e 1/1 to 1/2
NetIron(config)# write memory
NetIron(config)# trunk deploy
```

See the *Foundry Switch and Router Installation and Basic Configuration Guide* for more information on setting up trunk groups.

To configure a VLL like the one in Figure 5.3, enter commands such as the following:

```
NetIron(config)# router mpls
NetIron(config-mpls)# vll trunk_group 40000
NetIron(config-mpls-vll)# vll-peer 10.10.10.10
NetIron(config-mpls-vll)# untagged e 1/1
```

See the *Foundry NetIron Service Provider Configuration and Management Guide* for more information on configuring VLLs.

Notes:

- If you first create a trunk group, then configure a VLL, the port you specify as the VLL endpoint must also be the port you specified as the primary port of the trunk group.
- If you first configure a VLL, then create a trunk group, the primary port of the trunk group must be the same port you specified as the endpoint of the VLL.
- If you configure a port to be the endpoint of a VLL, then make that port the primary port of a trunk group, the VLL will use all the ports of the trunk group.
- If you later delete the trunk group from the configuration, traffic for the VLL is then handled solely by the port that had served as the primary port in the trunk group.
- If you specified a tagged endpoint for the VLL, all of the ports in the trunk group must be tagged.
- Untagged traffic received from any port in the trunk group is forwarded to the VLL. Tagged traffic is matched to its VLAN.
- When used as the endpoint of a VLL, a trunk group cannot have an odd number of ports. Only trunk groups with 2, 4, or 8 ports are supported.
- Only server trunk groups can be used as the endpoint of a VLL. Switch trunk groups are not supported for this purpose.
- If a port in the trunk group goes down, traffic is redistributed across the remaining ports. This may cause traffic flowing on one port to be moved to another port, even though neither is the port that went down.

Displaying MPLS VLL Information

You can display the following information about the MPLS VLL configuration on the Foundry device:

- Information about individual MPLS VLLs configured on the device
- Information about LDP sessions between VLL peers

Displaying Information About MPLS VLLs

To display information about MPLS VLLs:

```
NetIron# show mpls vll brief
Name              VC-ID   Vll-peer       End-point             State Tunnel-LSP
vll-1             1       2.2.2.2        untag e 2/1           UP    lsp_1
vll-2             2       3.3.3.3        tag vlan 200  e 2/2   DOWN  --
vll-3             3       --             undefined             DOWN  --
foundry-sj-to-sf  40000   192.168.2.100  untagged e 2/1        UP    lsp_1
```

Syntax: show mpls vll brief

For each MPLS VLL on the device, the following information is displayed:

Table 5.1: Output from the show mpls vll brief command

This Field...	Displays...
Name	The configured name of the VLL.
VC-ID	The user-configurable ID as defined in draft-martini-l2circuit-trans-mpls-07.txt.
Vll-peer	The remote PE router. This should be the same as the LSP destination for the LSPs that the VLL is transported over.
End-point	How packets are forwarded once they reach the egress LER. This can be one of the following: “untagged <portnum>” – Forward the packet out the specified port as untagged. “tag VLAN <vlan_id> <portnum>” – Tag the packet with the specified VLAN ID and forward the packet out the specified port. “undefined” – An endpoint has not been configured for this VLL
State	The current state of the VLL. This can be either UP or DOWN. Data can be forwarded over the VLL only when the state is UP.
Tunnel-LSP	The name of the RSVP-signalled LSP that has been selected to carry the VLL traffic through the MPLS domain

To display detailed information about the VLLs configured on the device:

```
NetIron# show mpls vll detail
VLL foundry-sj-to-sf VC-ID 40000
  State: UP
  Vll-peer:         2.2.2.2         End-point:        untagged  e 2/1
  Local label:      --              Remote label:     --
  Local group-id:   0               Remote group-id:  1
  COS:              --              Tunnel LSP:       lsp_1 (tnl2)
VLL vll-2 VC-ID 2
  State: DOWN - no tunnel LSP to vll-peer
  Vll-peer:         3.3.3.3         End-point:        tagged  vlan 200   e 2/2
  Local label:      --              Remote label:     --
  Local group-id:   0               Remote group-id:  --
  COS:              5               Tunnel LSP:       --
VLL vll-3 VC-ID 3
  State: DOWN - configuration incomplete
  Vll-peer:         --              End-point:        undefined
  Local label:      --              Remote label:     --
  Local group-id:   0               Remote group-id:  --
  COS:              7               Tunnel LSP:       --
VLL foundry-sj-to-sf VC-ID 40000
  State: UP
  Vll-peer:         192.168.2.100   End-point:        untagged e 2/1
  Local label:      --              Remote label:     --
  Local group-id:   0               Remote group-id:  --
  COS:              5               Tunnel LSP:       lsp_1 (tnl2)
```

Syntax: show mpls vll detail | <vll-name>

For each configured VLL, the command displays the following information:

Table 5.2: Output from the show mpls vll detail command

This Field...	Displays...
State:	The current state of the VLL. This can be one of the following: “UP”: VLL is operational – packets can flow. “DOWN - configuration incomplete”: A required configuration statement is missing. “DOWN - endpoint port to CE is down”: The physical endpoint port that should be connected to the Customer Edge device is down due to a link outage or is administratively disabled. "DOWN - no tunnel LSP to vll-peer": Cannot find a working LSP. "DOWN - no LDP session to vll-peer": LDP session is not yet ready. "DOWN - waiting for VC label binding from vll-peer": The device has advertised its VC label binding to the VLL peer, but has not yet received the peer’s VC label binding.
Vll-peer	The remote PE router. This should be the same as the LSP destination for the LSPs that the VLL is transported over.

Table 5.2: Output from the show mpls vll detail command (Continued)

This Field...	Displays...
End-point	How packets are forwarded once they reach the egress LER. This can be one of the following: "untagged <portnum>" – Forward the packet out the specified port as untagged. "tag VLAN <vlan_id> <portnum>" – Tag the packet with the specified VLAN ID and forward the packet out the specified port. "undefined" – An endpoint has not been configured for this VLL.
Local label:	The VC label value locally allocated for this VLL. Packets forwarded from the VLL peer to this device are expected to contain this label. This is the label that is advertised to the VLL peer through LDP.
Remote label:	The VC label allocated by the VLL peer and advertised to this device through LDP. The device applies this label to outbound MPLS packets sent to the VLL peer.
Local group-id:	The VLL group-ID (defined in draft-martini-l2circuit-trans-mpls-07.txt) advertised to the VLL peer through LDP. In this release, this is always zero.
Remote group-id:	The VLL group-ID selected and advertised by the VLL Peer.
COS:	The optional COS setting for the VLL. If a COS value is set, the Foundry device will attempt to select a tunnel LSP that also has this COS value. The COS value can be between 0 – 7.
Tunnel LSP:	The name, as well as internal tunnel index number, of the tunnel LSP selected for the VLL.

Displaying LDP Information

To display information about the state of the LDP connection between the Foundry device and VLL peers.

```
NetIron# show mpls ldp target-peer
Peer-addr      State
192.168.2.100  Initialized
```

Syntax: show mpls ldp target-peer

For each VLL peer, the command displays the following information:

Table 5.3: Output from the show mpls ldp target-peer command

This Field...	Displays...
Peer-addr:	The IP addresses of VLL peers.

Table 5.3: Output from the show mpls ldp target-peer command (Continued)

This Field...	Displays...
State:	The state of the LDP session with the VLL peer. This can be one of the following: “Unknown”: LDP session establishment has not started for this peer, normally because no Hello messages have been received from the peer. In this situation, that peer will not show up in the output of the **show mpls ldp session** command. "Nonexistent", "Initialized", "OpenRec", "OpenSent", or "Operational": LDP session states, as defined in RFC 3036.

To display information about LDP sessions between the Foundry device and VLL peers:

```
NetIron# show mpls ldp session
Peer LDP Ident: 192.168.2.100:1, Local LDP Ident: 11.1.1.1:1
  Active: no, State: Operational
  TCP connection: 11.1.1.1:646--22.2.2.2:9001, State: ESTABLISHED
  Addresses bound to peer LDP Ident:
    1.1.1.2
    10.1.1.2
    20.1.1.2
    22.2.2.2
```

Syntax: show mpls ldp session

For each established LDP session, the command displays the following information:

Table 5.4: Output from the show mpls ldp session command

This Field...	Displays...
Peer LDP Ident:	The VLL peer’s LDP identifier, consisting of the LSR ID and label space ID.
Local LDP Ident:	The Foundry device’s LDP identifier.
Active:	Whether this LSR is playing an active role in session establishment.
State:	The LDP session state, as defined in RFC 3036. This can be "Nonexistent", "Initialized", "OpenRec", "OpenSent", or "Operational".
TCP connection, state:	The TCP local/remote IP address, port and state.
Addresses bound to peer LDP Ident:	IP addresses carried in the VLL peer’s LDP Address messages.

Sample MPLS VLL Configuration

Figure 5.4 depicts a sample VLL configuration.

Figure 5.4 MPLS VLL Configuration

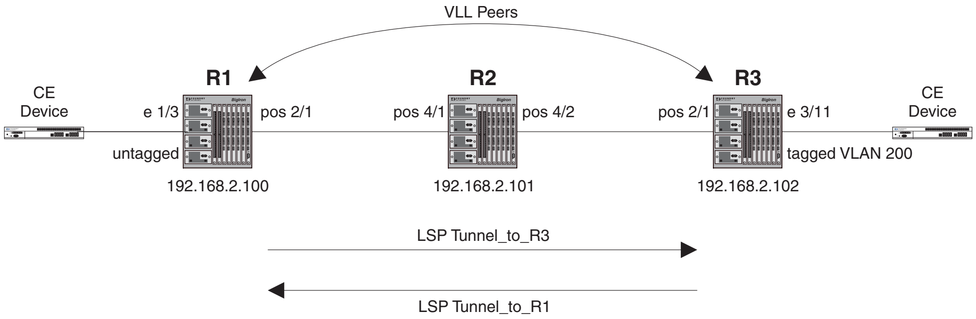

In this example, routers R1 and R3 are Provider Edge (PE) routers configured as VLL peers. R1 and R3 have established an LDP session to exchange VLL label information. When the LDP session is established, each router advertises its locally assigned VC label and VC ID to its VLL peer.

RSVP-signalled (tunnel) LSPs have been established in each direction between the two routers. When the CE device forwards a Layer 2 packet to R1, the router assigns the packet to an RSVP-signalled LSP whose destination is R3. R1 encapsulates the packet as an MPLS packet, adding a tunnel label and the VC label advertised to the router by R3. The MPLS packet is then forwarded over the outbound interface indicated by the tunnel label to the next hop in the LSP.

When the MPLS packet reaches R2, the penultimate LSR in the tunnel LSP, R2 pops the tunnel label, leaving the packet with only the VC label, then forwards the packet to R3.

R3 examines the VC label in the packet. On R3, the VC label is mapped to the user-specified endpoint for the VLL. In this example, the endpoint consists of VLAN ID 200 and interface 3/11. R3 then pops the VC label, tags the Layer 2 packet with VLAN 200, then forwards the packet out interface 3/11.

In the opposite direction, R3 assigns traffic received from the CE device to an RSVP-signalled LSP destined for R1, pushes tunnel and VC labels onto the packets, and forwards them to the next hop in LSP. When the packets reach R1, the router pops the VC label and forwards the Layer 2 packets out the interface indicated by the VLL endpoint. In this example, the endpoint consists of interface 1/3, so the packets are forwarded untagged out interface 1/3 to the CE device.

Router R1

The following commands configure Router R1 in Figure 5.4:

```
NetIron(config)# ip router-id 192.168.2.100
NetIron(config)# router ospf
NetIron(config-ospf-router)# area 0
NetIron(config-ospf-router)# exit

NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# traffic-engineering ospf
NetIron(config-mpls-policy)# exit

NetIron(config-mpls)# lsp Tunnel_To_R3
NetIron(config-mpls-lsp)# to 192.168.2.102
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# vll VLL_to_R3 40000
NetIron(config-mpls-vll)# vll-peer 192.168.2.102
NetIron(config-mpls-vll)# untagged e 1/3
NetIron(config-mpls-vll)# exit
NetIron(config-mpls)# interface loopback 1
NetIron(config-lbif-1)# port-name Generic All-Purpose Loopback
NetIron(config-lbif-1)# ip address 192.168.2.100/32
NetIron(config-lbif-1)# ip ospf area 0
NetIron(config-lbif-1)# exit

NetIron(config-mpls)# interface e 1/3
NetIron(config-if-e100-1/3)# port-name VLL_endpoint
NetIron(config-if-e100-1/3)# enable
NetIron(config-if-e100-1/3)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# port-name Connection_to_R2
NetIron(config-posif-2/1)# enable
NetIron(config-posif-2/1)# ip address 192.168.37.1/30
NetIron(config-posif-2/1)# ip ospf area 0
NetIron(config-posif-2/1)# exit
```

Router R2

The following commands configure Router R2 in Figure 5.4:

```
NetIron(config)# ip router-id 192.168.2.101
NetIron(config)# router ospf
NetIron(config-ospf-router)# area 0
NetIron(config-ospf-router)# exit

NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos

NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# traffic-engineering ospf
NetIron(config-mpls-policy)# exit

NetIron(config-mpls)# interface pos 4/1
NetIron(config-posif-4/1)# enable
NetIron(config-posif-4/1)# ip address 192.168.40.1/30
NetIron(config-posif-4/1)# ip ospf area 0
NetIron(config-posif-4/1)# exit

NetIron(config-mpls)# interface pos 4/2
NetIron(config-posif-4/2)# enable
NetIron(config-posif-4/2)# ip address 192.168.40.9/30
NetIron(config-posif-4/2)# ip ospf area 0
NetIron(config-posif-4/2)# exit

NetIron(config-mpls)# interface loopback 1
NetIron(config-lbif-1)# port-name Generic All-Purpose Loopback
NetIron(config-lbif-1)# ip address 192.168.2.101/32
NetIron(config-lbif-1)# ip ospf area 0
NetIron(config-lbif-1)# exit
```

Router R3

The following commands configure Router R3 in Figure 5.4:

```
NetIron(config)# ip router-id 192.168.2.102
NetIron(config)# router ospf
NetIron(config-ospf-router)# area 0
NetIron(config-ospf-router)# exit

NetIron(config)# router mpls
NetIron(config-mpls)# mpls-interface all-pos
NetIron(config-mpls)# policy
NetIron(config-mpls-policy)# traffic-engineering ospf
NetIron(config-mpls-policy)# exit

NetIron(config-mpls)# lsp Tunnel_To_R1
NetIron(config-mpls-lsp)# to 192.168.2.100
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)# exit

NetIron(config-mpls)# vll VLL_to_R1 40000
NetIron(config-mpls-vll)# vll-peer 192.168.2.100
NetIron(config-mpls-vll)# vlan 200
NetIron(config-mpls-vll-vlan)# tagged e 3/11
NetIron(config-mpls-vll-vlan)# exit
NetIron(config-mpls-vll)# exit

NetIron(config-mpls)# interface loopback 1
NetIron(config-lbif-1)# port-name Generic All-Purpose Loopback
NetIron(config-lbif-1)# ip address 192.168.2.102/32
NetIron(config-lbif-1)# ip ospf area 0
NetIron(config-lbif-1)# exit

NetIron(config-mpls)# interface e 3/11
NetIron(config-if-e100-3/11)# port-name VLL_endpoint
NetIron(config-if-e100-3/11)# enable
NetIron(config-if-e100-3/11)# exit

NetIron(config-mpls)# interface pos 2/1
NetIron(config-posif-2/1)# port-name Connection_to_R2
NetIron(config-posif-2/1)# enable
NetIron(config-posif-2/1)# ip address 192.168.41.1/30
NetIron(config-posif-2/1)# ip ospf area 0
NetIron(config-posif-2/1)# exit
```

Chapter 6
Configuring IS-IS

The ***Intermediate System to Intermediate System (IS-IS)*** protocol is a link-state Interior Gateway Protocol (IGP) that routers (intermediate systems) can use to exchange routes within a single routing domain. IS-IS is based on the International Standard for Organization/International Electrotechnical Commission (ISO/IEC) Open Systems Interconnect (OSI) networking model, and describes communication within the Networking layer of the model.

IS-IS is supported in the Service Provider software images on the following Foundry Networks product:

- NetIron Internet Backbone router

NOTE: References in this chapter to Layer 3 Switches also apply to the NetIron Internet Backbone router.

NOTE: The current software release does not support configuration or management of IS-IS using the Web management interface or IronView.

Overview

The following sections describe the Foundry implementation of the IS-IS protocol.

Specifications

The Foundry implementation of IS-IS is based on the following specifications and draft specifications:

- ISO/IEC 10589 – "Information Technology – Telecommunication and information exchange between systems – Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connection less-mode Network Service (ISO 8473)", 1992
- ISO/IEC 8473 – "Information processing systems – Data Communications – Protocols for providing the connectionless-mode network service", 1988
- ISO/IEC 9542 – "Information Technology – Telecommunication and information exchange between systems – End system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connection less-mode Network Service (ISO 8473)", 1988
- RFC 1195 – "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", 1990.
- RFC 1377 – "The PPP OSI Network Layer Control Protocol (OSINLCP)", 1992.
- RFC 2763 – "Dynamic Host Name Exchange Mechanism for IS-IS", 2000.
- RFC 2966 – "Domain-wide Prefix Distribution with Two-Level IS-IS", 2000
- Internet Draft – "IS-IS extensions for Traffic Engineering", 2000.

NOTE: Foundry supports the portions of this draft that describe the Extended IP reachability TLV (TLV type 135) and the extended Intermediate System (IS) reachability TLV (TLV type 22) to provide support for wide metrics.

NOTE: The Layer 3 Switch does not support routing of Connectionless-Mode Network Protocol (CLNP) packets. The Layer 3 Switch uses IS-IS for TCP/IP only.

Relationship to IP Route Table

The IS-IS protocol has the same relationship to the Layer 3 Switch's IP route table that OSPF has to the table. The protocol sends the best IS-IS path to a given destination to the CPU for comparison to the best paths from other protocols to the same destination. The CPU selects the path with the lowest administrative distance and places that path in the IP route table.

- If the path provided by IS-IS has the lowest administrative distance, then the CPU places that IS-IS path in the IP route table.
- If a path to the same destination supplied by another protocol has a lower administrative distance, the CPU installs the other protocol's path in the IP route table instead.

The ***administrative distance*** is a protocol-independent value from 1 – 255. Each path sent to the CPU, regardless of the source of the path (IS-IS, OSPF, static IP route, and so on) has an administrative distance.

Each route source has a default administrative distance. The default administrative distance for IS-IS is 115.

You can change the administrative distance for IS-IS and other routes sources.

For more information, see the following:

- The "IP Packet Flow Through a Layer 3 Switch" section in the "Configuring IP" chapter of the *Foundry Enterprise Configuration and Management Guide*.
- "Changing the Administrative Distance for IS-IS" on page 6-21

Intermediate Systems and End Systems

IS-IS uses the following categories to describe devices within an IS-IS routing domain (similar to an OSPF Autonomous System):

- ***Intermediate System (IS)*** – A device capable of forwarding packets from one device to another within the domain. In Internet Protocol (IP) terminology, an IS is a router. (Foundry routers are called "Layer 3 Switches".)
- ***End System (ES)*** – A device capable of generating or receiving packets within the domain. In IP terminology, an ES is an end node or IP host.

When you configure IS-IS on a Foundry Layer 3 Switch, the device is an IS.

Figure 6.1 shows an example of an IS-IS network.

Figure 6.1 An IS-IS network contains Intermediate Systems (ISs) and host systems

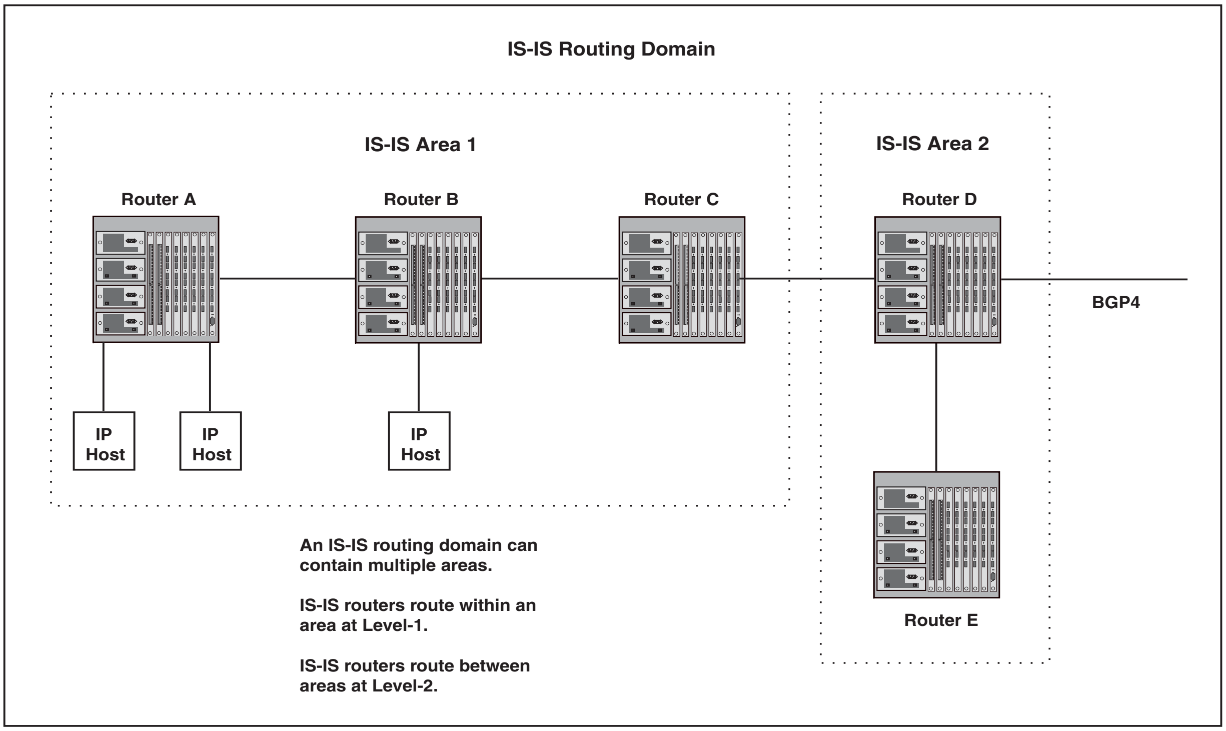
IS-IS Routing Domain
IS-IS Area 1
IS-IS Area 2
Router A
Router B
Router C
Router D
Router E
BGP4
IP Host
IP Host
IP Host
An IS-IS routing domain can contain multiple areas.
IS-IS routers route within an area at Level-1.
IS-IS routers route between areas at Level-2.

NOTE: Since the Foundry implementation of IS-IS does not route OSI traffic but instead routes IP traffic, IP hosts are shown instead of ESs.

The other basic IS-IS concepts illustrated in this figure are explained in the following sections.

Domain and Areas

IS-IS is an IGP, and thus applies only to routes within a single routing domain. However, you can configure multiple areas within a domain. A Foundry Layer 3 Switch can be a member of one area for each Network Entity Title (NET) you configure on the Layer 3 Switch. The NET contains the area ID for the area the NET is in.

In Figure 6.1, Routers A, B, and C are in area 1. Routers D and E are in area 2. All the routers are in the same domain.

Level-1 Routing and Level-2 Routing

You can configure an IS-IS router such as a Foundry Layer 3 Switch to perform one or both of the following levels of IS-IS routing[1]:

- Level-1 – A Level-1 router routes traffic only within the area the router is in. To forward traffic to another area, the Level-1 router sends the traffic to its nearest Level-2 router.
- Level-2 – A Level-2 router routes traffic between areas within a domain.

In Figure 6.1 on page 6-3, Routers A and B are Level-1 ISs only. Routers C and D are Level-1 ISs and Level-2 ISs. Router E is a Level-1 ISs only.

[1] The ISO/IEC specifications use the spelling “routeing”, but this document uses the spelling “routing” to remain consistent with other Foundry documentation.

Neighbors and Adjacencies

A Layer 3 Switch configured for IS-IS forms an ***adjacency*** with each of the IS-IS devices to which it is directly connected. An adjacency is a two-way direct link (a link without router hops) over which the two devices can exchange IS-IS routes and other protocol-related information. The link is sometimes called a "circuit". The devices with which the Layer 3 Switch forms adjacencies are its IS-IS ***neighbors***, which are other ISs.

A circuit can be a broadcast circuit or a point-to-point circuit. Foundry IS-IS interfaces are configured by default for broadcast circuits, but you can change the circuit type on an interface to point-to-point. Each end of an IS-IS adjacency must use the same circuit type.

In Figure 6.1 on page 6-3, Router A has an IS-IS adjacency with Router B. Likewise, Router B has an IS-IS adjacency with Router A and Router C.

Designated IS

A ***Designated IS*** is an IS-IS router that is responsible for gathering and distributing link state information to other Level-1 or Level-2 ISs within the same broadcast network (LAN). The Level-1 and Level-2 Designated ISs within a broadcast network are independent, although the same Layer 3 Switch can be a Level-1 Designated IS and a Level-2 Designated IS at the same time.

The Designated IS is elected based on the priority of each IS in the broadcast network. When an IS becomes operational, it sends a Level-1 or Level-2 Hello PDU to advertise itself to other ISs. If the IS is configured to be both a Level-1 and a Level-2 IS, the IS sends a separate advertisement for each level.

- The Level-1 IS that has the highest priority becomes the Level-1 Designated IS for the broadcast network.
- The Level-2 IS that has the highest priority becomes the Level-2 Designated IS for the broadcast network.

If the Designated IS becomes unavailable (for example, is rebooted), the IS with the next highest priority becomes the new IS. If two or more ISs have the highest priority, the IS with the highest MAC address becomes the Designated IS.

The priority is an interface parameter. Each interface that is enabled for IS-IS can have a different priority.

Figure 6.2 shows an example of the results of Designated IS elections. For simplicity, this example shows four of the five routers in Figure 6.1 on page 6-3, withe same domain and areas.

Figure 6.2 Each broadcast network has a Level-1 Designated IS and a Level-2 Designated IS

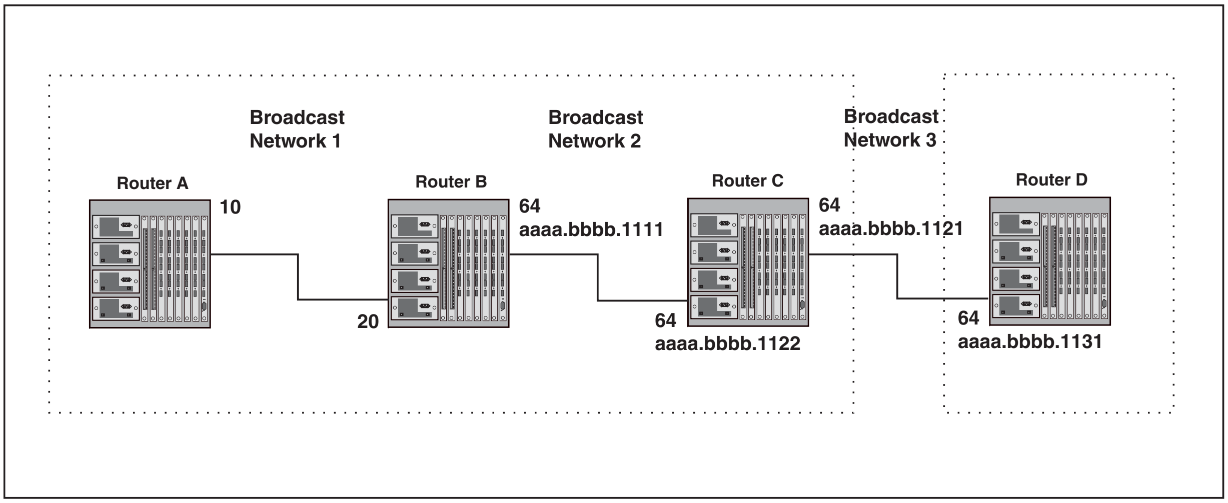

Designated IS election has the following results in this network topology:

- Router B is the Level-1 Designated IS for broadcast network 1
- Router C is the Level-1 Designated IS for broadcast network 2
- Router D is the Level-2 Designated IS for broadcast network 3

In this example, the IS-IS priorities for the IS-IS interfaces in broadcast network 1 have been changed by an administrator. The priorities for the interfaces in the other broadcast networks are still set to the default (64). When there is a tie, IS-IS selects the interface with the highest MAC address.

Broadcast Pseudonode

In a broadcast network, the Designated IS maintains and distributes link state information to other ISs by maintaining a ***pseudonode***. A pseudonode is a logical host representing all the Level-1 or Level-2 links among the ISs in a broadcast network. Level-1 and Level-2 have separate pseudonodes, although the same device can be the pseudonode for Level-1 and Level-2.

Route Calculation and Selection

The Designated IS uses a ***Shortest Path First (SPF)*** algorithm to calculate paths to destination ISs and ESs. The SPF algorithm uses Link State PDUs (LSPDUs) received from other ISs as input, and creates the paths as output.

After calculating the paths, the Designated IS then selects the best paths and places them in the IS-IS route table. The Designated IS uses the following process to select the best paths:

1. Prefer the Level-1 path over the Level-2 path.
2. If there is no Level-1 path, prefer the internal Level-2 path over the external Level-2 path.
3. If there is still more than one path, prefer the path with the lowest metric.
4. If there is more than one path with the lowest metric, load share among the paths.

After selecting the best path to a destination, the software places the path in the IS-IS route table.

IS-IS Parameters and Defaults

This section describes the IS-IS parameters you can configure on a global or individual interface basis, and describes when the parameter changes take effect.

When Parameter Changes Take Effect

Some parameter changes take effect immediately while others do not take full effect until you disable, then re-enable route redistribution. None of the IS-IS parameters require a software reload to places changes into effect.

Immediately

The following parameter changes take effect immediately:

- Enable or disable IS-IS.
- Add a NET.
- Set the overload bit on or off.
- Change the priority (used during election of the Designated IS).
- Change the supported IS-IS level (Level-1 or Level-2).
- Change subdomain or area authentication.
- Disable or re-enable adjacency formation.
- Disable or re-enable padding of Hello PDUs.
- Change the metric style (narrow or wide).
- Change the maximum number of load-sharing paths.
- Change the CSNP and PSNP interval.
- Change the maximum LSP lifetime.
- Change the maximum LSP refresh interval.

- Change the LSP general interval.
- Change the SPF timer.
- Enable default route origination.
- Change the administrative distance.
- Change the default metric.
- Add a summary address.
- Enable or disable route redistribution.
- Disable or enable display of hostnames (RFC 2763 mapping of the IS-IS system ID to the hostname).
- Clear neighbor sessions, routes, PDU statistics, or error statistics.

After Disabling and Re-Enabling Redistribution

The following parameter change takes effect only after you disable and then re-enable redistribution:

- Change the default metric.
- Add, change, or negate route redistribution parameters.

Global Parameters

Table 6.1 lists the global IS-IS parameters for Layer 3 Switches.

Table 6.1: Global IS-IS Parameters

Parameter	Description	Default	See page...
Required Parameters The following parameters must be configured before the Layer 3 Switch can perform IS-IS routing.			
IS-IS state	The state of the protocol **Note**: You also must enable the protocol locally on individual interfaces to activate the protocol on those interfaces.	Disabled	6-11
Network Entity Title (NET)	The Layer 3 Switch's IS-IS identity. You can configure more than one NET on the device.	None configured	6-11
Neighbor Parameters The following parameters apply to the Layer 3 Switch's relationships with its IS-IS neighbors.			
Overload bit	A bit in the header of an LSP (link-state PDU) that indicates whether the Layer 3 Switch has sufficient resources to perform Level-1 or Level-2 IS-IS routing. The bit can be on (1) or off (0): • On (1) – The Layer 3 Switch does not have sufficient resources for the indicated level. • Off (0) – The Layer 3 Switch has sufficient resources for the indicated level. You can use the overload bit to administratively shut down IS-IS routing, either to make configuration changes or to allow the protocol to fully come up following a software reload.	Off (0)	6-13

Table 6.1: Global IS-IS Parameters (Continued)

Parameter	Description	Default	See page...
IS-IS level	The level of IS-IS routing enabled on the Layer 3 Switch. The level can be set to one of the following: • Level-1 and Level-2 • Level-1 only • Level-2 only	Level-1 and Level-2	6-14
Domain authentication	A password the Layer 3 Switch uses to authenticate Level-2 LSPDUs that the Layer 3 Switch sends or receives.	None configured	6-15
Area authentication	A password the Layer 3 Switch uses to authenticate Level-1 LSPDUs that the Layer 3 Switch sends or receives.	None configured	6-15
Interface authentication	A password the Layer 3 Switch uses to authenticate Hello PDUs that the Layer 3 Switch sends or receives on the specified interface.	None configured	6-16
Hello padding	Extra bytes in a hello PDU sent by an IS-IS router that indicate the maximum length of IS-IS PDU the router can accept. The padding makes the Hello PDU the same length as the Maximum Transmission Unit (MTU).	Enabled	6-16
Metric style	The length of the metric field in Link State PDUs. The style can be one of the following: • Narrow – one byte • Wide – four bytes The wide metric style enables you to use extensions to the IS-IS for Traffic Engineering (TE).	Narrow	6-17
Route Parameters The following parameters apply to IS-IS routes.			
Maximum load sharing paths	The maximum number of equal-cost paths across which the Layer 3 Switch is allowed to distribute traffic.	4	6-18
Sequence Numbers PDU interval	How often a Layer 3 Switch acting as the Designated IS sends Level-1 or Level-2 Complete Sequence Numbers Protocol Data Units (CSNPs) or Partial Sequence Numbers Protocol Data Units (PSNPs). You can specify a value from 0 – 65535 seconds. The interval applies to Level-1 and Level-2 CSNPs.	10 seconds	6-19
Maximum LSP lifetime	The maximum number of seconds an unrefreshed LSP can remain in the Layer 3 Switch’s LSP database. The maximum LSP lifetime can be from 1 – 65535 seconds.	1200 seconds (20 minutes)	6-19

Table 6.1: Global IS-IS Parameters (Continued)

Parameter	Description	Default	See page...
LSP refresh interval	The maximum number of seconds the Layer 3 Switch waits between sending updated LSPs to its IS-IS neighbors. The interval can be from 1 – 65535 seconds.	900 seconds (15 minutes)	6-19
LSP general interval	The minimum number of seconds the Layer 3 Switch waits between sending updated LSPs to its IS-IS neighbors. The interval can be from 1 – 120 seconds.	10 seconds	6-20
SPF timer	How often the Layer 3 Switch recalculates the Shortest Path First (SPF) tree of its IS-IS links following a change in topology or the link state database. You can specify a value from 1 – 120 seconds.	5 seconds	6-20
Default route advertisement	Whether the Layer 3 Switch advertises a default route to its IS-IS neighbors. When default route advertisement is enabled, the feature applies to Level-2 only. However, you can use route maps to also readvertise the default route on Level-1.	Disabled	6-21
Administrative distance	A protocol-independent metric that the Layer 3 Switch uses when evaluating equal-cost paths from different routing protocols to the same destination. When comparing otherwise equal routes from different protocols, the Layer 3 Switch selects the route with the lower administrative distance. You can specify a value from 1 – 255.	115	6-21
Summary addresses	A single address that the Layer 3 Switch uses in place of multiple addresses for routing. Summary addresses help enhance performance and reduce the size of the Link State database.	None configured	6-22
Default redistribution metric	The cost that the Layer 3 Switch assigns to a route redistributed from another source into IS-IS, if the route does not already have a valid metric value. The default metric can be a value from 1 – 65535.	10	6-23

Table 6.1: Global IS-IS Parameters (Continued)

Parameter	Description	Default	See page...
Route redistribution	Whether the Layer 3 Switch can exchange routes between IS-IS and other routing protocols. You can configure redistribution of routes from the following sources into IS-IS: • IP static routes • IP routes to directly connected devices • RIP routes • OSPF routes • BGP4 routes You also can redistribute routes between IS-IS Level-1 and Level-2.	Disabled	6-22
Monitoring Parameters The following parameters apply to IS-IS information display and monitoring.			
Display of the IS-IS hostname	Display of the mapping between the Layer 3 Switch's IS-IS system ID and its hostname.	Enabled	6-25
Logging of adjacency changes	The Layer 3 Switch can generate Syslog entries and SNMP traps to indicate changes in the status of an adjacency with another IS.	Disabled	6-26

Interface Parameters

Table 6.2 lists the IS-IS interface parameters for Layer 3 Switches.

Table 6.2: IS-IS Interface Parameters

Parameter	Description	Default	See page...
Required Parameters The following parameters must be configured before the Layer 3 Switch can perform IS-IS routing.			
IS-IS state	The state of the protocol **Note**: You also must enable the protocol globally.	Disabled	6-11
Neighbor Parameters The following parameters apply to the Layer 3 Switch's relationships with its IS-IS neighbors.			

Table 6.2: IS-IS Interface Parameters (Continued)

Parameter	Description	Default	See page...
Priority	The priority of this Layer 3 Switch relative to other IS-IS routers for election as the Designated IS. Level-1 and Level-2 each have their own Designated ISs. You can specify a priority from 0 – 255. A higher numeric value means a higher priority. **Note**: Priority settings apply to both Level-1 and Level-2 by default, but you can specify Level-1 or Level-2 when setting the priority.	64	6-14
IS-IS type	Overrides the global setting. See Table 6.1 on page 6-6.	Level-1 and Level-2	6-14
Area authentication	A password the Layer 3 Switch uses to authenticate Level-1 LSPDUs that the Layer 3 Switch sends or receives.	None configured	6-15
Adjacency formation	Whether the Layer 3 Switch is enabled to form a Level-2 adjacency (peer relationship) with the IS at the other end of the link on this interface. This parameter can have one of the following values: • Active – The interface can form adjacencies with ISs at the other end of the link. • Passive – The interface is advertised into IS-IS, but the interface does not send advertisements. **Note**: This parameter does not affect advertisement of the interface into the IS-IS area. Advertisement is always enabled.	Loopback interfaces – passive All other interfaces – Active	6-16
Hello padding	Overrides the global setting. See Table 6.1 on page 6-6.	Enabled	6-16
Hello interval	How often the Layer 3 Switch sends hello messages to its IS-IS neighbors. You can specify from 1 – 65535 seconds. You also can set the interval for Level-1 only, Level-2 only, or both.	10 seconds, for Level-1 and Level-2	6-17
Hello multiplier	The number by which the Layer 3 Switch multiplies the hello interval to obtain the hold time for Level-1 and Level-2 IS-to-IS hello PDUs. You can set the multiplier to a value from 1 – 2147483647.	3	6-18
Route Parameters The following parameters apply to IS-IS routes.			
Metric	The cost that the Layer 3 Switch adds to routes originated on the interface or when calculating routes.	10	6-20

Basic Configuration Tasks

You must enable IS-IS globally, then enable the protocol locally on the interfaces attached to ISs or ESs. After enabling IS-IS globally and on individual interfaces, you must configure a Network Entity Title (NET) on each interface. The Layer 3 Switch does not begin exchanging information with IS-IS neighbors until you enable IS-IS on the interfaces attached to the ISs and ESs.

The following sections describe the configuration tasks that are required to use a Foundry Layer 3 Switch as an IS-IS router.

Enabling IS-IS and Configuring the NET

IS-IS is disabled by default. To use the protocol, you must do the following:

- Globally enable IS-IS.
- Globally configure at least one Network Entity Title (NET). The NET is the Layer 3 Switch's network interface with IS-IS. You can configure up to three NETs on the Layer 3 Switch.
- Enable IS-IS on the individual interfaces that are attached to ISs or ESs.

You must enable the protocol both globally and on individual interfaces. The NETs are global and thus apply to all IS-IS interfaces on the Layer 3 Switch.

USING THE CLI

To globally enable IS-IS on the Layer 3 Switch, enter commands such as the following at the global CONFIG level of the CLI:

```
NetIron(config)# router isis
ISIS: Please configure NET!
```

The command in this example globally enables IS-IS and changes the CLI to the IS-IS configuration level. If you have not already configured a NET for the IS-IS, the message shown in this example is displayed. To configure a NET, enter a command such as the following:

```
NetIron(config-isis-router)# net 49.2211.aaaa.bbbb.cccc.00
```

The command in this example configures a NET that has the area ID 49.2211, the system ID aaaa.bbbb.cccc (the Layer 3 Switch's base MAC address), and SEL value 00.

In addition to enabling IS-IS globally, you also must enable the protocol on the individual interfaces connected to ISs or ESs. To enable IS-IS locally on specific interfaces, enter commands such as the following:

```
NetIron(config)# interface ethernet 1/1
NetIron(config-if-1/1)# ip router isis
NetIron(config-if-1/1)# interface ethernet 1/2
NetIron(config-if-1/2)# ip router isis
```

These commands enable IS-IS on ports 1/1 and 1/2. The NET configured above (at the IS-IS configuration level) applies to both interfaces.

NOTE: If you have not configured a NET, the software displays the message "ISIS: Please configure NET!" and changes the CLI to the IS-IS configuration level.

Syntax: [no] router isis

Syntax: [no] net <area-id>.<system-id>.<sel>

Syntax: [no] ip router isis

The <area-id> parameter specifies the area and has the format xx or xx.xxxx. For example, 49 and 49.2211 are valid area IDs.

The <system-id> parameter specifies the Layer 3 Switch's unique IS-IS router ID and has the format xxxx.xxxx.xxxx. You can specify any value for the system ID. A common practice is to use the device's base MAC address as the system ID. The base MAC address is also the MAC address of port 1/1. To determine the base

MAC address, enter the following command at any level of the CLI: **show interfaces brief**. The base MAC address is listed in the first row of information, in the MAC column.

You must use the same system ID in all the NETs on the Layer 3 Switch.

NOTE: The parameter descriptions above are the recommended values for the NET. However, the CLI accepts any value that fits within the following lengths and formats:

xx.xxxx.xxxx.xxxx.00 – minimum length of NET

xx.xxxx.xxxx.xxxx.xxxx.xxxx.xxxx.xxxx.xxxx.xxxx.00 – maximum length of NET

The <sel> parameter specifies the NSAP Selector (SEL). This value must always be 00 (two zeros). The value 00 indicates that this address is an NET.

Configuring IS-IS on an ATM Interface

IS-IS is supported on point-to-point and point-to-multipoint ATM interfaces.

Point-To-Multipoint

A point-to-multipoint interface means the interface can have multiple point-to-point connections to other devices in the same ATM cloud. With this feature, one ATM interface is sufficient to establish adjacencies with multiple remote routing devices.

IS-IS internally treats the point-to-multipoint interface as a broadcast circuit. All the protocol functions and features that pertain to IS-IS broadcast interfaces also apply to IS-IS interfaces configured on ATM point-to-multipoint interfaces. LAN IS-IS Hello PDUs are sent to discover neighbors.

A fully meshed ATM cloud is not required for the support of IS-IS point-to-multipoint. However, a device that intends to establish an adjacency with a router needs to configure its interface as point-to-multipoint.

To configure an IS-IS interface on an ATM point-to-multipoint interface, enter commands such as the following:

```
BigIron(config)# interface atm 3/1.4 multipoint
BigIron(config-if-3/1.4)# atm pvc 3 34 ubr ip 30.0.0.2
BigIron(config-if-3/1.4)# atm pvc 3 37 ubr ip 30.0.0.1
BigIron(config-if-3/1.4)# ip address 30.0.0.3 255.0.0.0
BigIron(config-if-3/1.4)# ip router isis
BigIron(config-if-3/1.4)# exit
BigIron(config)# ip router isis
ISIS: Please configure NET!
BigIron(config-isis-router)# net 49.2211.aaaa.bbbb.cccc.00
BigIron(config-isis-router)# exit
BigIron(config)# interface ethernet 3/1
BigIron(config-if-3/1)# ip router isis
```

Point-To-Point

IS-IS support over ATM point-to-point interface is similar to that on POS interface. Point-to-point neighbor adjacencies are formed and point-to-point IS-IS Hello PDUs are sent to discover neighbors.

To configure an IS-IS interface on an ATM point-to-point interface, enter commands such as the following:

```
BigIron(config)# interface atm 3/1.4
BigIron(config-if-3/1.4)# atm pvc 3 34 ubr
BigIron(config-if-3/1.4)# ip address 30.0.0.3 255.0.0.0
BigIron(config-if-3/1.4)# exit
BigIron(config)# ip router isis
ISIS: Please configure NET!
BigIron(config-isis-router)# net 49.2211.aaaa.bbbb.cccc.00
BigIron(config-isis-router)# exit
BigIron(config)# interface ethernet 3/1
BigIron(config-if-3/1)# ip router isis
```

Optional Configuration Tasks

Most of the IS-IS parameters have defaults. After you perform the configuration tasks in “Basic Configuration Tasks” on page 6-11, the Layer 3 Switch performs IS-IS routing using the default settings. You can change parameter settings for the following:

- Neighbors and adjacencies – see “Configuring Neighbor and Adjacency Parameters”
- Routes – see “Configuring Route Parameters” on page 6-18

You also can change the following display and monitoring parameters:

- Display of the Layer 3 Switch’s IS-IS hostname – see “Disabling or Re-enabling Display of the Layer 3 Switch Hostname” on page 6-25
- Logging of adjacency changes – see “Logging Adjacency Changes” on page 6-26

Configuring Neighbor and Adjacency Parameters

This section describes how to change the following parameters related to neighbors and forming adjacencies with them:

- Overload Bit
- Priority for Designated IS election
- IS-IS level (Level-1 or Level-2)
- Authentication
- Adjacency formation on specific IS-IS interfaces
- Padding of Hello messages
- Metric style (wide or narrow)
- Changing the Hello Interval
- Changing the Hello Multiplier

Setting the Overload Bit

If an IS’s resources are overloaded, preventing the IS from properly performing IS-IS routing, the IS can inform other ISs of this condition by setting the overload bit in LSPDUs sent to other ISs from 0 (off) to 1 (on).

When an IS is overloaded, other ISs will not use the overloaded IS to forward traffic. An IS can be in the overload state for Level-1, Level-2, or both.

- If an IS is in the overload state for Level-1, other Level-1 ISs stop using the overloaded IS to forward Level-1 traffic. However, the IS can still forward Level-2 traffic, if applicable.
- If an IS is in the overload state for Level-2, other Level-2 ISs stop using the overloaded IS to forward Level-2 traffic. However, the IS can still forward Level-1 traffic, if applicable.
- If an IS is in the overload state for both levels, the IS cannot forward traffic at either level.

By default, the Layer 3 Switch automatically sets the overload bit to 1 (on) in its LSPDUs to other ISs if an overload condition occurs.

You can set the overload bit on to administratively shut down IS-IS without disabling the protocol. Setting the overload bit on is useful when you want to make configuration changes without removing the Layer 3 Switch from the network.

In addition, you can configure the Layer 3 Switch to set the overload bit on for a specific number of seconds during startup, to allow IS-IS to become fully active before the device begins IS-IS routing. By default, there is no delay (0 seconds).

USING THE CLI

To immediately set the overload bit on, enter the following command:

```
NetIron(config-isis-router)# set-overload-bit
```

This command administratively shuts down IS-IS by configuring the Layer 3 Switch to immediately set the overload bit to 1 (on) in all LSPs sent to other ISs.

To configure the Layer 3 Switch to temporarily set the overload bit on after a software reload, enter a command such as the following:

```
NetIron(config-isis-router)# set-overload-bit on-startup 5
```

This command configures the Layer 3 Switch to set the overload bit on in all IS-IS LSPs sent to other ISs during the first five seconds following a successful software reload. After the five seconds expire, the Layer 3 Switch stops setting the overload bit on, and instead starts setting the overload bit off.

Syntax: [no] set-overload-bit [on-startup <secs>]

The **on-startup** <secs> parameter specifies the number of seconds following a reload to set the overload bit on. You can specify 0 or a number from 5 – 86400 (24 hours). The default is 0, which means the Layer 3 Switch starts performing IS-IS routing immediately following a successful software reload.

Setting the Priority for Designated IS Election

The priority of an IS-IS interface determines the priority of the interface for being elected as a Designated IS. Level-1 has a Designated IS and Level-2 has a Designated IS. The Level-1 and Level-2 Designated ISs are independent, although the same device can become both the Level-1 Designated IS and the Level-2 Designated IS.

By default, the Level-1 and Level-2 priority is 64. You can configure an interface's priority to a value from 0 – 255. You can configure the same priority for both Level-1 and Level-2 or you can configure a different priority for each level. In case of a tie (if two or more devices have the highest priority within a given level), the device with the highest MAC address becomes the Designated IS for that level.

NOTE: You can set the IS-IS priority on an individual interface basis only. You cannot set the priority globally.

To set the IS-IS priority on an interface, use either of the following methods.

USING THE CLI

To set the IS-IS priority on an interface, enter commands such as the following:

```
NetIron(config-isis-router)# interface ethernet 1/1
NetIron(config-if-1/1)# isis priority 128
```

This command sets the IS-IS priority on port 1/1 to 128. Since the command does not specify Level-1 or Level-2, the new priority setting applies to both IS-IS levels.

Syntax: [no] isis priority <num> [level-1 | level-2-only]

The <num> parameter specifies the priority and can be from 0 – 255. A higher numeric value means a higher priority. The default is 64.

The **level-1 | level-2-only** parameter applies the priority to Level-1 or Level-2 only. By default, the priority is applied to both levels.

Changing the IS-IS Level

By default, a Foundry Layer 3 Switch can operate as both a Level-1 and IS-IS Level-2 router. You can change the IS-IS type globally or on an individual interface to be Level-1 only or Level-2 only. You also can reset the type to both Level-1 and Level-2.

NOTE: If you change the IS-IS type on an individual interface, the type you specify must also be specified globally. For example, if you globally set the type to Level-2 only, you cannot set the type on an individual interface to Level-1. The software accepts the setting but the setting does not take effect.

USING THE CLI

To globally change the type of IS-IS packets supported on the device from Level-1 and Level-2 to Level-1 only, enter the following command:

```
NetIron(config-isis-router)# is-type level-1
```

Syntax: [no] is-type level-1 | level-1-2 | level-2-only

The **level-1 | level-1-2 | level-2-only** parameter specifies the IS-IS type. If you want to re-enable support for both IS-IS types, re-enter the command you entered to change the IS-IS type, and use “no” in front of the command. For example, to reverse the command shown above an re-enable support for both IS-IS types, enter the following command:

```
NetIron(config-isis-router)# no is-type level-1
```

If you want to change the IS-IS type on a specific interface only, enter commands such as the following:

```
NetIron(config-isis-router)# interface ethernet 1/1
NetIron(config-if-1/1)# isis circuit-type level-1
```

These commands change the CLI to the interface configuration level, then change the IS-IS type supported for the IS-IS circuit on that interface.

Syntax: [no] isis circuit-type level-1 | level-1-2 | level-2

Configuring Authentication

By default, the Layer 3 Switch does not authenticate packets sent to or received from ESs or other ISs. You can configure the following types of passwords for IS-IS.

Table 6.3: IS-IS Passwords

Password Type	Scope	Where Used	Default
Domain	Level-2	Level-2 LSPDU	None configured
Area	Level-1	Level-2 LSPDU	None configured
Interface	Level-1 and Level-2	Hello PDU	None configured

If you configure a password, the Layer 3 Switch checks for the password in IS-IS packets received by the device and includes the password in packets sent by the device. For example, the Layer 3 Switch checks all Level-2 LSPDUs received by the device for the domain password you configure, and includes the password in all Level-2 PDUs sent by the device.

Configuring a Domain Password

To configure an IS-IS domain password, use the following method.

USING THE CLI

To configure an IS-IS domain password, enter a command such as the following:

```
NetIron(config-isis-router)# domain-password domain-1
```

This command configures the Foundry device to use the password “domain-1” to authenticate Level-2 LSPDUs.

Syntax: [no] domain-password <string>

The <string> parameter specifies the password. You can enter an alphanumeric string up to 80 characters long. The password can contain blank spaces. If you use a blank space in the password, you must use quotation marks (“ “) around the entire password; for example, **domain-password “domain 1”**.

Configuring an Area Password

To configure an IS-IS area password, use the following method.

USING THE CLI

To configure an IS-IS area password, enter a command such as the following:

```
NetIron(config-isis-router)# area-password area-51
```

This command configures the Foundry device to use the password “area-51” to authenticate Level-1 LSPDUs.

Syntax: [no] area-password <string>

The <string> parameter specifies the password. You can enter an alphanumeric string up to 80 characters long. The password can contain blank spaces. If you use a blank space in the password, you must use quotation marks (“ “) around the entire password; for example, **area-password “area 51”**.

Configuring an Interface Password

To configure an IS-IS interface password, use the following method.

USING THE CLI

To configure an IS-IS password on an interface, enter commands such as the following:

```
NetIron(config)# interface ethernet 1/1
NetIron(config-if-1/1)# isis password int1
```

This command configures the Foundry device to use the password “int1” to authenticate Hello PDUs sent or received on port 1/1.

Syntax: [no] isis password <string>

The <string> parameter specifies the password. You can enter an alphanumeric string up to 80 characters long. The password can contain blank spaces. If you use a blank space in the password, you must use quotation marks (“ “) around the entire password; for example, **isis password “interface 1”**.

Disabling or Re-Enabling Formation of Adjacencies

When you enable IS-IS on any type of interface except a loopback interface, the interface also is enabled to send advertisements and form an adjacency with an IS at the other end of the link by default. Adjacency formation and advertisements are disabled by default on loopback interfaces.

You can enable or disable adjacency formation and advertisements on an interface.

NOTE: The Foundry device advertises an IS-IS interface to its area regardless of whether adjacency formation is enabled.

USING THE CLI

To disable IS-IS adjacency formation on an interface, enter commands such as the following:

```
NetIron(config-isis-router)# interface ethernet 1/1
NetIron(config-if-1/1)# isis passive
```

This command disables IS-IS adjacency formation on port 1/1. The device still advertises this IS-IS interface into the area, but does not allow the port to form an adjacency with the IS at the other end of the link.

Syntax: [no] isis passive

Disabling or Re-Enabling Hello Padding

By default, the Layer 3 Switch adds extra data to the end of a hello packet to make the packet the same size as the maximum length of PDU the Layer 3 Switch supports.

The padding applies to the following types of hello packets:

- ES hello (ESH PDU)
- IS hello (ISH PDU)
- IS to IS hello (IIH PDU)

The padding consists of arbitrarily valued octets. A padded hello PDU indicates the largest PDU that the Layer 3 Switch can receive. Other ISs that receive a padded hello PDU from the Layer 3 Switch can therefore ensure that the IS-IS PDUs they send the Layer 3 Switch. Similarly, if the Layer 3 Switch receives a padded hello PDU from a neighbor IS, the Layer 3 Switch knows the maximum size PDU that the Layer 3 Switch can send to the neighbor.

When padding is enabled, the maximum length of a Hello PDU sent by the Layer 3 Switch is 1514 bytes.

If you need to disable padding, you can do so globally or on individual interfaces. Generally, you do not need to disable padding unless a link is experiencing slow performance, for example due to point-to-point interoperability issues. If you enable or disable padding on an interface, the interface setting overrides the global setting.

By default, disabling or re-enabling padding affects hello PDUs sent on point-to-point circuits and to an IS-IS broadcast address. You can specify an option to enable or disable the padding for point-to-point or broadcast PDUs.

USING THE CLI

To globally disable padding of IS-IS hello PDUs, enter the following command:

```
NetIron(config-isis-router)# no hello padding
```

This command disables all hello PDU padding on the Layer 3 Switch. To re-enable padding, enter the following command:

```
NetIron(config-isis-router)# hello padding
```

To disable padding on a specific interface only, enter commands such as the following:

```
NetIron(config-isis-router)# interface ethernet 1/1
NetIron(config-if-1/1)# hello padding
```

Syntax: [no] hello padding [point-to-point]

The **point-to-point** parameter disables or re-enables the padding only for point-to-point connections.

Changing the Metric Style

The metric style specifies the Types, Lengths, and Values (TLVs) an IS-IS LSP can have. The TLVs specify the types of data, the maximum length of the data, and the valid values for the data. One of the types of data the TLVs control is a route’s default-metric. By default, the Layer 3 Switch uses the standard IS-IS TLVs, which allows metric values from 1 – 63. The default metric style is called “narrow”. You can increase the range of metric values supported by the Layer 3 Switch by changing the metric style to wide. The wide metric style allows metric values from 1 – 16777215.

You can change the metric style for one or both levels (Level-1 and Level-2).

USING THE CLI

To change the metric style to wide, enter the following command:

```
NetIron(config-isis-router)# metric-style wide
```

This command changes the metric style for both Level-1 and Level-2.

Syntax: [no] metric-style wide [level-1 | level-2-only]

The **level-1 | level-1-2 | level-2-only** parameter specifies the level(s) to which the change applies.

Changing the Hello Interval

The hello interval controls how often the Layer 3 Switch sends hello messages to its IS-IS neighbors. The default interval is 10 seconds for Level-1 and Level-2. You can change the hello interval for one or both levels to a value from 1 – 65535 seconds.

USING THE CLI

To change the hello interval, enter a command such as the following:

```
NetIron(config-isis-router)# hello-interval 20
```

This command changes the hello interval to 20 seconds. By default, the change applies to both Level-1 and Level-2.

Syntax: [no] hello-interval <num> [level-1 | level-2-only]

The <num> parameter specifies the interval, and can be from 1 – 65535 seconds. The default is 10 seconds.

The **level-1 | level-2-only** parameter applies the change to only the level you specify. If you do not use this parameter, the change applies to both levels.

Changing the Hello Multiplier

The hello multiplier is the number by which the Layer 3 Switch multiplies the hello interval to obtain the hold time for Level-1 and Level-2 IS-to-IS hello PDUs. The default multiplier is 3. You can set the multiplier to a value from 1 – 2147483647.

USING THE CLI

To change the hello multiplier, enter a command such as the following:

```
NetIron(config-isis-router)# hello-multiplier 50
```

This command changes the hello interval to 50. By default, the change applies to both Level-1 and Level-2.

Syntax: [no] hello-multiplier <num> [level-1 | level-2-only]

The <num> parameter specifies the multiplier, and can be from 1 – 2147483647. The default is 3.

The **level-1 | level-2-only** parameter applies the change to only the level you specify. If you do not use this parameter, the change applies to both levels.

Configuring Route Parameters

This section describes how to change the following parameters related to LSPDUs and routes:

- Maximum number of load sharing paths
- CSNP and PSNP Interval
- Maximum LSP Lifetime
- LSP Refresh Interval
- LSP General Interval
- SPF Timer
- Metric (cost) the Layer 3 Switch adds to IS-IS routes before advertising them
- Advertisement of the default route
- IS-IS administrative distance
- Summary addresses
- Route redistribution

Changing the Maximum Number of Load Sharing Paths

When the Layer 3 Switch has multiple IS-IS equal-cost paths to the same destination, the Layer 3 Switch can load share among the paths. For example, if the device has three IS-IS paths to the same destination and each path has the same default metric, the Layer 3 Switch alternates among the three paths when forwarding traffic.

By default, IS-IS load sharing is enabled for four paths. You can change the number of paths to an amount from 1 – 8. If you change the number of paths to 1, the Layer 3 Switch does not load share route paths learned from IS-IS.

NOTE: IS-IS load sharing requires IP load sharing to be enabled for at least the number of paths you want to use for IS-IS load sharing. IP load sharing is enabled for four paths by default. The IP load sharing settings affect all routing protocols that support load sharing, including IS-IS, OSPF and BGP4.

USING THE CLI

To change the maximum number of IS-IS load sharing paths, enter a command such as the following:

```
NetIron(config-isis-router)# maximum-paths 6
```

Syntax: [no] maximum-paths <num>

The <num> parameter specifies the maximum number of load sharing paths and can be from 1 – 8. The default is 4.

Changing the Sequence Numbers PDU Interval

A ***Complete Sequence Numbers PDU (CSNP)*** is a complete list of the LSPs in the Designated IS' link state database. The CSNP contains a list of all the LSPs in the database, as well as other information that helps IS neighbors determine whether their LSP databases are in sync with one another. The Designated IS sends CSNPs to the broadcast interface. Level-1 and Level-2 each have their own Designated IS.

A ***Partial Sequence Numbers PDU (PSNP)*** is a partial list of LSPs. ISs other than the Designated IS (that is, the non-Designated ISs) send PSNPs to the broadcast interface.

The CSNP interval specifies how often the Designated IS sends a CSNP to the broadcast interface. Likewise, the PSNP interval specifies how often other ISs (non-Designated ISs) send a PSNP to the broadcast interface. (The PSNP interval also applies to ISs on a point-to-point network.)

The interval you can configure on the Layer 3 Switch applies to both Level-1 and Level-2 CSNPs and PSNPs. The default interval is 10 seconds. You can set the interval to a value from 0 – 65535 seconds.

USING THE CLI

To change the interval, enter a command such as the following:

```
NetIron(config-isis-router)# csnp-interval 15
```

Syntax: [no] csnp-interval <secs>

The <secs> parameter specifies the interval and can be from 0 – 65535 seconds. The default is 10 seconds.

NOTE: Although the command name is **csnp-interval**, the interval also applies to PSNPs.

Changing the Maximum LSP Lifetime

The maximum LSP lifetime is the maximum number of seconds an unrefreshed LSP can remain in the Layer 3 Switch's LSP database. The maximum LSP lifetime can be from 1 – 65535 seconds. The default is 1200 seconds (20 minutes).

USING THE CLI

To change the maximum LSP lifetime to 2400 seconds, enter a command such as the following:

```
NetIron(config-isis-router)# max-lsp-lifetime 2400
```

Syntax: [no] max-lsp-lifetime <secs>

The <secs> parameter specifies the maximum LSP lifetime and can be from 1 – 65535 seconds. The default is 1200 seconds (20 minutes).

Changing the LSP Refresh Interval

The LSP refresh interval is the maximum number of seconds the Layer 3 Switch waits between sending updated LSPs to its IS-IS neighbors. The interval can be from 1 – 65535 seconds. The default is 900 seconds.

USING THE CLI

To change the LSP refresh interval to 20000 seconds, enter a command such as the following:

```
NetIron(config-isis-router)# lsp-refresh-interval 20000
```

Syntax: [no] lsp-refresh-interval <secs>

The <secs> parameter specifies the maximum refresh interval and can be from 1 – 65535 seconds. The default is 900 seconds (15 minutes).

Changing the LSP General Interval

The LSP general interval is the minimum number of seconds the Layer 3 Switch waits between sending updated LSPs to its IS-IS neighbors. The interval can be from 1 – 120 seconds. The default is 10 seconds.

USING THE CLI

To change the LSP general interval to 45 seconds, enter a command such as the following:

```
NetIron(config-isis-router)# lsp-gen-interval 45
```

Syntax: [no] lsp-gen-interval <secs>

The <secs> parameter specifies the minimum refresh interval and can be from 1 – 120 seconds. The default is 10 seconds.

Changing the SPF Timer

Every IS maintains a Shortest Path First (SPF) tree, which is a representation of the states of each of the IS’s links to ESs and other ISs. If the IS is both a Level-1 and Level-2 IS, it maintains separate SPF trees for each level.

To ensure that the SPF tree remains current, the IS updates the tree at regular intervals following a change in network topology or the link state database. By default, the Foundry Layer 3 Switch recalculates its IS-IS tree every five seconds following a change. You can change the SPF timer to a value from 1 – 120 seconds.

USING THE CLI

To change the SPF interval, enter a command such as the following:

```
NetIron(config-isis-router)# spf-interval 30
```

Syntax: [no] spf-interval <secs>

The <secs> parameter specifies the interval and can be from 1 – 120 seconds. The default is 5 seconds.

Changing the Metric Added to Advertised Routes

When the Layer 3 Switch originates an IS-IS route or calculates a route, the Layer 3 Switch adds a metric (cost) to the route. Each IS-IS interface has a separate metric value. The default is 10.

The Layer 3 Switch applies the interface-level metric to routes originated on the interface and also when calculating routes. The Layer 3 Switch does not apply the metric to link-state information that the Layer 3 Switch receives from one IS and floods to other ISs.

The default interface metric is 10. You can change the metric on an individual interface to a value in one of the following ranges:

- 1 – 63 for the narrow metric style (the default metric style)
- 1 – 16777215 for the wide metric style

NOTE: If the metric value you want to use is higher than 63 but you have not changed the metric style to wide, change the metric style first, then set the metric. To change the metric style, see “Changing the Metric Style” on page 6-17. The IS-IS neighbors that will receive the advertisements also must be enabled to receive wide metrics.

To change the IS-IS metric on an interface, use the following CLI method.

USING THE CLI

```
NetIron(config-isis-router)# interface ethernet 1/1
NetIron(config-if-1/1)# isis metric
```

Syntax: [no] isis metric <num>

The <num> parameter specifies the metric. The range of values you can specify depends on the metric style. You can specify 1 – 63 for the narrow metric style or 1 – 16777215 for the wide metric style. The default in either case is 10.

Enabling Advertisement of a Default Route

By default, the Foundry device does not generate or advertise a default route to its neighboring ISs. This is true even if the device's IP route table contains a default route. You can enable the Foundry device to advertise a default route to all neighboring ISs using one of the following methods. By default, the feature originates the default route at Level-2 only. However, you can apply a route map to originate the default route to Level-1 only or at both Level-1 and Level-2.

NOTE: This feature requires the presence of a default route in the IP route table.

USING THE CLI

To enable the Layer 3 Switch to advertise a default route, enter the following command:

```
NetIron(config-isis-router)# default-information-originate
```

This command enables the device to advertise a default route into the IS-IS area to which the device is attached.

Syntax: [no] default-information-originate [route-map <name>]

The **route-map** <name> parameter allows you to specify the level to which to advertise the default route. You can specify one of the following:

- Advertise to Level-1 ISs only.
- Advertise to Level-2 ISs only.
- Advertise to Level-1 and Level-2 ISs.

NOTE: The route map must be configured before you can use the route map as a parameter with the **default-information-originate** command.

To use a route map to specify the level on which you want the Layer 3 Switch to advertise a default route, change the CLI to the global CONFIG level, then enter commands such as the following:

```
NetIron(config)# route-map DarkAngel permit 1
NetIron(config-routemap DarkAngel)# set level level-1
NetIron(config-routemap DarkAngel)# router isis
NetIron(config-isis-router)# default-information-originate route-map DarkAngel
```

These commands configure a route map to set the default advertisement level to Level-1 only.

Syntax: [no] route-map <map-name> permit | deny <num>

Syntax: [no] set level level-1 | level-1-2 | level-2

For this use of a route map, use the **permit** option and do not specify a match statement. Specify a set statement to set the level to one of the following:

- **level-1** – Level-1 only
- **level-1-2** – Level-1 and Level-2
- **level-2** – Level-2 only

Changing the Administrative Distance for IS-IS

When the Layer 3 Switch has paths from multiple routing protocols to the same destination, the Layer 3 Switch compares the administrative distances of the paths and selects the path with the lowest administrative distance to place in the IP route table.

For example, if the Layer 3 Switch has a path from RIP, from OSPF, and a path from IS-IS to the same destination, and all the paths are using their protocols' default administrative distances, the Layer 3 Switch selects the OSPF path, because that path has a lower administrative distance than the RIP and IS-IS paths.

Here are the default administrative distances on the Foundry Layer 3 Switch:

- Directly connected – 0 (this value is not configurable)

- Static – 1 (applies to all static routes, including default routes)
- EBGP – 20
- OSPF – 110
- IS-IS – 115
- RIP – 120
- IBGP – 200
- Local BGP – 200
- Unknown – 255 (the router will not use this route)

Lower administrative distances are preferred over higher distances. For example, if the Layer 3 Switch receives routes for the same network from IS-IS and from RIP, the Layer 3 Switch will prefer the IS-IS route by default.

USING THE CLI

To change the administrative distance for IS-IS routes, enter a command such as the following:

```
NetIron(config-isis-router)# distance 100
```

Syntax: [no] distance <num>

This command changes the administrative distance for all IS-IS routes to 100.

The <num> parameter specifies the administrative distance. You can specify a value from 1 – 255. The default for IS-IS is 115.

Configuring Summary Addresses

You can configure summary addresses to aggregate IS-IS route information. Summary addresses can enhance performance by reducing the size of the Link State database, reducing the amount of data the Layer 3 Switch needs to send to its neighbors, and reducing the CPU cycles used for IS-IS.

When you configure a summary address, the address applies only to Level-2 routes by default. You can specify Level-1 only, Level-2 only, or Level-1 and Level-2 when you configure the address.

USING THE CLI

To configure a summary address, enter a command such as the following:

```
NetIron(config-isis-router)# summary-address 192.168.0.0 255.255.0.0
```

This command configures a summary address for all Level-2 IS-IS route destinations between 192.168.1.0 – 192.168.255.255.

Syntax: [no] summary-address <ip-addr> <ip-mask> [level-1 | level-1-2 | level-2-only]

The <ip-addr> <ip-mask> parameters specify the aggregate address. The mask indicates the significant bits in the address. Ones are significant, and zeros allow any value. In the command example above, the mask 255.255.0.0 matches on all addresses that begin with 192.168 and contain any values for the final two octets.

The **level-1 | level-1-2 | level-2-only** parameter specifies the route types to which the aggregate route applies. The default is **level-2-only**.

Redistributing Routes

The Layer 3 Switch can redistribute routes from the following route sources into IS-IS:

- BGP4
- RIP
- OSPF
- Static
- Directly connected

The Layer 3 Switch also can redistribute Level-1 IS-IS routes into Level-2 IS-IS routes, and Level-2 routes into Level-1 routes.

Route redistribution from other sources into IS-IS is disabled by default. When you enable redistribution, the Layer 3 Switch redistributes routes only into Level-2 by default. You can specify Level-1 only, Level-2 only, or Level-1 and Level-2 when you enable redistribution.

The Layer 3 Switch automatically redistributes Level-1 routes into Level-2 routes. Thus, you do not need to enable this type of redistribution. You also can enable redistribution of Level-2 routes into Level-1 routes.

The Layer 3 Switch attempts to use the redistributed route's metric as the route's IS-IS metric. For example, is an OSPF route has an OSPF cost of 20, the Layer 3 Switch uses 20 as the route's IS-IS metric. The Layer 3 Switch uses the redistributed route's metric as it IS-IS metric unless either of the following occurs:

- If the route does not a have a valid metric, the Layer 3 Switch assigns the default metric value to the route.
- If the metric value is valid but also is higher than the maximum metric value supported by the Layer 3 Switch, the Layer 3 Switch uses the maximum value that is supported as the route's cost. For example, if the metric style is narrow (1 – 63), but a redistributed route's metric is 99, the Layer 3 Switch gives the route an IS-IS metric of 63.

Changing the Default Redistribution Metric

When IS-IS redistributes a route from another route source (such as OSPF, BGP4, or an IP static route) into IS-IS, IS-IS uses the route's metric value as its IS-IS metric. However, if the route does not have a valid metric, IS-IS instead applies a default metric to the route. The default value for the default metric is 10. You can change the default metric to a value from 1 – 65535.

NOTE: By default, the Layer 3 Switch supports metric values from 1 – 63. To use a higher metric value, you must change the metric style to "wide". See "Changing the Metric Style" on page 6-17.

NOTE: The Foundry implementation of IS-IS does not support the optional metric types Delay, Expense, or Error.

USING THE CLI

To change the default metric, enter a command such as the following:

```
NetIron(config-isis-router)# default-metric 20
```

Syntax: [no] default-metric <num>

The <num> parameter specifies the default metric. You can specify a value from 1 – 65535. The default is 10.

Redistributing Static IP Routes into IS-IS

To redistribute static routes from the IP route table into IS-IS routes, enter the following command:

```
NetIron(config-isis-router)# redistribute static
```

This command configures the Layer 3 Switch to redistribute all IP static routes into Level-2 IS-IS routes.

Syntax: [no] redistribute static [level-1 | level-1-2 | level-2 | metric <num> |
metric-type external | internal |
route-map <name>]

The **level-1**, **level-1-2**, and **level-2** parameters restrict redistribution to the specified IS-IS level.

The **metric** <num> parameter restricts the redistribution to only those routes that have the metric you specify.

The **metric-type external | internal** parameter restricts redistribution to one of the following:

- **external** – The metric value is not comparable to an IS-IS internal metric and is always higher than the IS-IS internal metric.
- **internal** – The metric value is comparable to metric values used by IS-IS. This is the default.

The **route-map** <name> parameter restricts distribution to those routes that match the specified route map. The route map must already be configured before you use the route map name with the **redistribute** command. For example, to configure a route map that redistributes only the static IP routes to the destination network 2.4.69.x, enter commands such as the following:

```
NetIron(config)# access-list 101 permit ip any 2.4.69.0 255.255.255.0
NetIron(config)# route-map Alba permit 1
NetIron(config-routemap DarkAngel)# match ip address 101
NetIron(config-routemap DarkAngel)# router isis
NetIron(config-isis-router)# redistribute static route-map Alba
```

For information about the ACL and route map syntax, see the following:

- The "IP Access Control Lists (ACLs)" chapter in the *Foundry Enterprise Configuration and Management Guide*.
- The "Defining Route Maps" section in the "Configuring BGP4" chapter of the *Foundry Enterprise Configuration and Management Guide*.

Redistributing Directly Connected IP Routes into IS-IS

To redistribute directly connected IP routes into IS-IS routes, enter the following command:

```
NetIron(config-isis-router)# redistribute connected
```

This command configures the Layer 3 Switch to redistribute all directly connected routes in the IP route table into Level-2 IS-IS.

Syntax: [no] redistribute connected [level-1 | level-1-2 | level-2 | metric <num> |
metric-type external | internal |
route-map <name>]

The parameters are the same as the parameters for the **redistribute static** command.

Redistributing RIP Routes into IS-IS

To redistribute RIP routes into IS-IS, enter the following command:

```
NetIron(config-isis-router)# redistribute rip
```

This command configures the Layer 3 Switch to redistribute all RIP routes into Level-2 IS-IS.

Syntax: [no] redistribute rip [level-1 | level-1-2 | level-2 | metric <num> |
metric-type external | internal |
route-map <name>]

The parameters are the same as the parameters for the **redistribute static** command.

Redistributing OSPF Routes into IS-IS

To redistribute OSPF routes into IS-IS, enter the following command:

```
NetIron(config-isis-router)# redistribute ospf
```

This command configures the Layer 3 Switch to redistribute all OSPF routes into Level-2 IS-IS.

Syntax: [no] redistribute ospf [level-1 | level-1-2 | level-2 |
match external1 | external2 | internal
metric <num> |
metric-type external | internal |
route-map <name>]

Most of the parameters are the same as the parameters for the **redistribute static** command. However, the **redistribute ospf** command also has the **match external1 | external2 | internal** parameter. This parameter specifies the OSPF route type you want to redistribute into IS-IS. By default, the **redistribute ospf** command redistributes only internal routes.

- **external1** – An OSPF type 1 external route.
- **external2** – An OSPF type 2 external route.

- **internal** – An internal route calculated by OSPF.

Redistributing BGP4 Routes into IS-IS

To redistribute BGP4 routes into IS-IS, enter the following command:

```
NetIron(config-isis-router)# redistribute bgp
```

This command configures the Layer 3 Switch to redistribute all its BGP4 routes into Level-2 IS-IS.

Syntax: [no] redistribute bgp [level-1 | level-1-2 | level-2 | metric <num> |
metric-type external | internal |
route-map <name>]

The parameters are the same as the parameters for the **redistribute static** command.

Redistributing IS-IS Routes Within IS-IS

In addition to redistributing routes from other route sources into IS-IS, you also can redistribute routes from one IS-IS level to the other. By default, the Layer 3 Switch redistributes routes from Level-1 into Level-2.

NOTE: The Layer 3 Switch automatically redistributes Level-1 routes into Level-2 routes, even if you do not enable redistribution.

USING THE CLI

To redistribute IS-IS routes from one level to the other, enter a command such as the following:

```
NetIron(config-isis-router)# redistribute isis level-2 into level-1
```

This command redistributes the Level-2 routes into Level-1. The Layer 3 Switch automatically redistributes Level-1 routes into Level-2.

Syntax: [no] redistribute isis level-1 into level-2 | level-2 into level-1

The **level-1 into level-2 | level-2 into level-1** parameter specifies the direction of the redistribution:

- **level-1 into level-2** – Redistributes Level-1 routes into Level-2. This is the default.
- **level-2 into level-1** – Redistributes Level-2 routes into Level-1.

Disabling or Re-enabling Display of the Layer 3 Switch Hostname

Foundry's implementation of IS-IS supports RFC 2763, which describes a mechanism for mapping IS-IS system IDs to the hostnames of the devices with those IDs. For example, if you set the hostname on the Layer 3 Switch to "IS-IS Router 1", the mapping feature uses this name instead of the Layer 3 Switch's IS-IS system ID in the output of the following commands:

- **show isis database**
- **show isis interface**
- **show isis neighbor**

The Layer 3 Switch's hostname is displayed in each CLI command prompt, for example:

```
IS-IS Router 1(config-isis-router)#
```

The name mapping feature is enabled by default. If you want to disable name mapping, enter the following command:

```
IS-IS Router 1(config-isis-router)# no hostname
```

Syntax: [no] hostname

To display the name mappings, enter the **show isis hostname** command. See "Displaying the Name Mappings" on page 6-27.

Logging Adjacency Changes

The Layer 3 Switch can generate a Syslog entry and an SNMP trap to indicate a change in the status of an adjacency with another IS. Logging of the adjacency changes is disabled by default. To enable or disable them, use either of the following methods.

To display the Syslog messages, see “Displaying IS-IS Syslog Messages” on page 6-28.

USING THE CLI

To enable logging of adjacency changes, enter the following command:

```
NetIron(config-isis-router)# log-adjacency-changes
```

Syntax: [no] log-adjacency-changes

To disable logging of adjacency changes, enter the following command:

```
NetIron(config-isis-router)# no log-adjacency-changes
```

Displaying IS-IS Information

You can display the following information:

- The active configuration (the IS-IS commands in the running-config) – see “Displaying the IS-IS Configuration in the Running-Config” on page 6-26
- Name mappings – “Displaying the Name Mappings” on page 6-27
- Neighbor information – “Displaying Neighbor Information” on page 6-27
- Neighbor adjacency changes – “Displaying IS-IS Syslog Messages” on page 6-28
- Interface information – “Displaying Interface Information” on page 6-30
- Route information – “Displaying Route Information” on page 6-33
- LSP database entries – “Displaying LSP Database Entries” on page 6-34
- Traffic statistics – “Displaying Traffic Statistics” on page 6-37
- Error statistics – “Displaying Error Statistics” on page 6-38

Displaying the IS-IS Configuration in the Running-Config

You can display the global IS-IS configuration commands that are in effect on the Layer 3 Switch using the following CLI method.

NOTE: The running-config does not list the default values. Only commands that change a setting or add configuration information are displayed. For information about IS-IS defaults, see “IS-IS Parameters and Defaults” on page 6-5.

USING THE CLI

To list the global IS-IS configuration commands in the Layer 3 Switch’s running-config, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis config

Current IS-IS configuration:
router isis
 net  20.00e0.5200.0001.00
end
```

The running-config shown in this example contains the command that enables IS-IS and a command that configures an NET.

To display the interface configuration information in the running-config, enter one of the following commands at any level of the CLI:

- **show running-config**
- **write terminal**

Syntax: show isis config

Displaying the Name Mappings

To display the mappings between IS-IS system IDs and the hostnames of the devices with those IDs, use the following CLI method.

USING THE CLI

To display the mappings, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis hostname
Total number of entries in IS-IS Hostname Table: 1
   System ID       Hostname            * = local IS
 * bbbb.cccc.dddd  BigIron
```

Syntax: show isis hostname

The table in this example contains one mapping, for this Layer 3 Switch. The Layer 3 Switch's IS-IS system ID is "bbbb.cccc.dddd" and its hostname is "BigIron". The display contains one entry for each IS that supports name mapping.

NOTE: Name mapping is enabled by default. When name mapping is enabled, the output of the **show isis database**, **show isis neighbor**, and **show isis routes** commands uses the host name instead of the system ID. To disable mapping so that these displays use the system ID instead, see "Disabling or Re-enabling Display of the Layer 3 Switch Hostname" on page 6-25.

Displaying Neighbor Information

To display information about the Layer 3 Switch's IS-IS neighbors, use either of the following methods.

USING THE CLI

To display IS-IS neighbor information, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis neighbor
Total number of IS-IS Neighbors: 2
System ID      Interface    SNPA           State Holdtime Type Pri StateChgeTime
00e0.52b5.7800 Ether2/4     00e0.52b5.7843 UP    10       ISL2 64  0   :0 :16:8
00e0.52b5.7800 Ether2/4     00e0.52b5.7843 UP    10       ISL1 64  0   :0 :16:8
```

Syntax: show isis neighbor

This display shows the following information.

Table 6.4: IS-IS Neighbor Information

This Field...	Displays...
Total number of IS-IS Neighbors	The number of ISs with which the Layer 3 Switch has formed IS-IS adjacencies.

Table 6.4: IS-IS Neighbor Information (Continued)

This Field...	Displays...
System ID	The System ID of the neighbor.
Interface	The Layer 3 Switch port or virtual interface attached to the neighbor.
SNPA	The Subnetwork Point of Attachment (SNPA), which is the MAC address of the Layer 3 Switch port or virtual interface attached to the neighbor.
State	The state of the adjacency with the neighbor. The state can be one of the following: • DOWN – The adjacency is down. • INIT – The adjacency is being established and is not up yet. • UP – The adjacency is up.
Holdtime	The time between transmission of IS-IS hello messages.
Type	The IS-IS type of the adjacency. The type can be one of the following: • ISL1 – Level-1 IS • ISL2 – Level-2 IS • PTP – Point-to-Point IS • ES – ES **Note**: The Layer 3 Switch forms a separate adjacency for each IS-IS type. Thus, if the Layer 3 Switch has both types of IS-IS adjacencies with the neighbor, the display contains a separate row of information for each adjacency.
Pri	The priority of this IS to be elected as the Designated IS in this broadcast network.
StateChgeTime	The amount of time that has passed since the adjacency last changed state.

Displaying IS-IS Syslog Messages

When logging is enabled, the Layer 3 Switch generates Syslog messages and SNMP traps for the following IS-IS events:

- Overload state (the Layer 3 Switch entering or leaving the overload state)
- Memory overrun (IS-IS is demanding more memory than is available)

You also can enable the Layer 3 Switch to generate Syslog messages and SNMP traps when an adjacency with a neighbor comes up or goes down. To enable logging of adjacency changes, see “Logging Adjacency Changes” on page 6-26.

To display Syslog entries, use the following method.

USING THE CLI

To display Syslog entries, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):
00d00h00m42s:N:BGP Peer 192.147.202.10 UP (ESTABLISHED)
00d00h00m18s:N:ISIS  L2 ADJACENCY  UP 1234.1234.1234 on circuit 2
00d00h00m08s:N:ISIS  L1 ADJACENCY  UP 1234.1234.1234 on circuit 2
00d00h00m08s:N:ISIS  L2 ADJACENCY  UP 0000.86de.5520 on circuit 1
00d00h00m00s:I:Warm start
```

The messages in this example indicate that the software has been reloaded (Warm start) and adjacencies between the Layer 3 Switch and three ISs have come up.

Syntax: show logging

Table 6.5 lists the IS-IS Syslog messages.

Table 6.5: IS-IS Syslog Messages

Message Level	Message	Explanation
Alert	ISIS MEMORY USE EXCEEDED	IS-IS is requesting more memory than is available.
Notification	ISIS L1 ADJACENCY DOWN <system-id> on circuit <circuit-id>	The Layer 3 Switch's adjacency with this Level-1 IS has gone down. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	ISIS L1 ADJACENCY UP <system-id> on circuit <circuit-id>	The Layer 3 Switch's adjacency with this Level-1 IS has come up. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	ISIS L2 ADJACENCY DOWN <system-id> on circuit <circuit-id>	The Layer 3 Switch's adjacency with this Level-2 IS has gone down. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.

Table 6.5: IS-IS Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	ISIS L2 ADJACENCY UP <system-id> on circuit <circuit-id>	The Layer 3 Switch's adjacency with this Level-2 IS has come up. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	ISIS ENTERED INTO OVERLOAD STATE	The Layer 3 Switch has set the overload bit to on (1), indicating that the Layer 3 Switch's IS-IS resources are overloaded.
Notification	ISIS EXITING FROM OVERLOAD STATE	The Layer 3 Switch has set the overload bit to off (0), indicating that the Layer 3 Switch's IS-IS resources are no longer overloaded.

Displaying Interface Information

To display information about the interfaces on which IS-IS is enabled, use either of the following methods.

USING THE CLI

To display information about the Layer 3 Switch's IS-IS interfaces, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis interface

Total number of IS-IS Interfaces: 2

Interface : 2/4     Local Circuit Number: 00000001
   Circuit Type : BCAST Circuit Mode : LEVEL-1-2
   Circuit State: UP Passive State: FALSE
   MTU : 1497
   Level-1 Metric: 10, Level-1 Priority: 64
   Level-1 Hello Interval: 10 Level-1 Hello Multiplier: 3
   Level-1 Designated IS: 00e0.52b5.7800.01-00  Level-1 DIS Changes: 4
   Level-2 Metric: 10, Priority: 64
   Level-2 Hello Interval: 10 Level-2 Hello Multiplier: 3
   Level-2 Designated IS: 00e0.52b5.7800.01-00, Level-2 DIS Changes: 5
   Next IS-IS LAN Level-1 Hello in 3 seconds
   Next IS-IS LAN Level-2 Hello in 8 seconds
   Number of active level-1 adjacencies: 1
   Number of active level-2 adjacencies: 1
   Circuit State Changes: 0 Circuit Adjacencies State Changes: 2
   Rejected Adjacencies: 0
   Circuit Authentication Fails: 0 Bad LSP: 0
   Control Messages Sent: 204 Control Messages Received: 1990
   IP Address and Subnet Mask:
     128.1.1.2          255.255.255.0
...
```

Syntax: show isis interface

This display shows the following information.

Table 6.6: IS-IS Interface Information

This Field...	Displays...
Total number of IS-IS interfaces	The number of interfaces on which IS-IS is enabled.
Interface	The port or virtual interface number to which the information listed below applies.
Local Circuit Number	The ID that the instance of IS-IS running on the interface applied to the circuit between this interface and the interface at the other end of the link.
Circuit Type	The type of IS-IS circuit running on the interface. The circuit type can be one of the following: • BCAST– broadcast • PTP – point-to-point
Circuit Mode	The IS-IS type in use on the circuit. The mode can be one of the following: • LEVEL-1 • LEVEL-2 • LEVEL-1-2
Circuit State	The state of the circuit, which can be one of the following: • DOWN • UP
Passive State	The state of the passive option, which determines whether the interface is allowed to form an IS-IS adjacency with the IS at the other end of the circuit. The state can be one of the following: • FALSE – The passive option is disabled. The interface can form an adjacency with the IS at the other end of the link. • TRUE – The passive option is enabled. The interface cannot form an adjacency, but can still advertise itself into the area.
MTU	The maximum length supported for IS-IS PDUs sent on this interface.
Level-1 Metric	The default-metric value that the Layer 3 Switch inserts in IS-IS Level-1 PDUs originated on this interface.
Level-1 Priority	The priority of this IS to be elected as the Designated IS for Level-1 in this broadcast network.
Level-1 Hello Interval	The number of seconds the software waits between sending Level-1 hello PDUs to the IS at the other end of the circuit.
Level-1 Hello Multiplier	The number by which the software multiplies the hello interval to calculate the hold time for Level-1 Hello messages received on the circuit.
Level-1 Designated IS	The NET of the Level-1 Designated IS.

Table 6.6: IS-IS Interface Information (Continued)

This Field...	Displays...
Level-1 DIS Changes	The number of times the NET of the Level-1 Designated IS has changed.
Level-2 Metric	The default-metric value that the Layer 3 Switch inserts in IS-IS Level-2 PDUs originated on this interface.
Level-2 Priority	The priority of this IS to be elected as the Designated IS for Level-2 in this broadcast network.
Level-2 Hello Interval	The number of seconds the software waits between sending Level-2 Hello messages to the IS at the other end of the circuit.
Level-2 Hello Multiplier	The number by which the software multiplies the hello interval to calculate the hold time for Level-2 LSPs received on the circuit.
Level-2 Designated IS	The NET of the Level-2 Designated IS.
Level-2 DIS Changes	The number of times the NET of the Level-2 Designated IS has changed.
Next IS-IS LAN Level-1 Hello	Number of seconds before next Level-1 Hello message will be transmitted by the Layer 3 Switch.
Next IS-IS LAN Level-2 Hello	Number of seconds before next Level-2 Hello message will be transmitted by the Layer 3 Switch.
Number of active Level-1 adjacencies	The number of ISs with which this interface has an active Level-1 adjacency.
Number of active Level-2 adjacencies	The number of ISs with which this interface has an active Level-2 adjacency.
Circuit State Changes	The number of times the state of the circuit has changed.
Circuit State Adjacencies Changes	The number of times an adjacency has started or ended on this circuit.
Rejected Adjacencies	The number of adjacency attempts by other ISs rejected by the Layer 3 Switch.
Circuit Authentication Fails	The number of times the Layer 3 Switch rejected a circuit because the authentication did not match the authentication configured on the Layer 3 Switch.
Bad LSP	The number of times the interface received a bad LSP from an IS at the other end of the circuit. The following conditions can cause an LSP to be bad: • Invalid checksum • Invalid length • Invalid lifetime value
Control Messages Sent	The number of IS-IS control PDUs sent on this interface.
Control Messages Received	The number of IS-IS control PDUs received on this interface.
IP Address and Subnet Mask	The IP address and sub-net mask configured on this interface.

Displaying Route Information

To display the routes in the Layer 3 Switch’s IS-IS route table, use either of the following methods.

USING THE CLI

To display information about the routes in the Layer 3 Switch’s IS-IS route table, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis routes

Total number of IS-IS routes: 26
Destination      Mask             Cost  Type Tag      Flags1    Flags2
50.50.15.0       255.255.255.0    11    L2   00000000 00000640  73010000
   Path: 1 Next Hop IP: 128.1.1.1         Interface: 2/4   Flags :84000003
50.50.18.0       255.255.255.0    11    L2   00000000 00000640  73010000
   Path: 1 Next Hop IP: 128.1.1.1         Interface: 2/4   Flags :84000003
50.50.21.0       255.255.255.0    11    L2   00000000 00000640  73010000
   Path: 1 Next Hop IP: 128.1.1.1         Interface: 2/4   Flags :84000003
```

Syntax: show isis routes

This display shows the following information.

Table 6.7: IS-IS Route Information

This Field...	Displays...
Total number of IS-IS routes	The total number of routes in the Layer 3 Switch’s IS-IS route table. The total includes Level-1 and Level-2 routes.
Destination	The IP destination of the route.
Mask	The sub-net mask for the destination address.
Cost	The IS-IS default metric for the route, which is the cost of using this route to reach the next-hop router to this destination.
Type	The route type, which can be one of the following: • L1 – Level-1 route • L2 – Level-2 route
Tag	The tag value associated with the route.
Flags1	Values used by Foundry technical support for troubleshooting.
Flags2	Values used by Foundry technical support for troubleshooting.
Path	The path number in the table. The IS-IS route table can contain multiple equal-cost paths to the same destination, in which case the paths are numbered consecutively from 1 up to 4. When IP load sharing is enabled, the Layer 3 Switch can load balance traffic to the destination across the multiple paths.
Next Hop IP	The IP address of the next-hop interface to the destination.
Interface	The Layer 3 Switch interface (port or virtual interface) attached to the next hop.
Flags	Values used by Foundry technical support for troubleshooting.

Displaying LSP Database Entries

Use the following methods to display summary or detailed information about the entries in the LSP database.

NOTE: The Layer 3 Switch maintains separate LSP databases for Level-1 LSPs and Level-2 LSPs.

Displaying Summary Information

To display summary information about the LSPs in the Layer 3 Switch's LSP databases, use either of the following methods.

USING THE CLI

To display summary information for all the LSPs in the Layer 3 Switch's LSP databases, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis database

IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
00e0.5200.0001.00-00* 0x00000009   0x027b        1082              0/0/1
00e0.52b5.7800.00-00  0x00000007   0x8631        1014              0/0/0
00e0.52b5.7800.01-00  0x00000006   0xcb17        1014              0/0/0

IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
00e0.5200.0001.00-00* 0x0000000a   0xc1da        1082              0/0/1
00e0.52b5.7800.00-00  0x00000005   0xf307        115               0/0/0
```

The command in this example shows information for the LSPS in the Layer 3 Switch's Level-1 and Level-2 LSP databases. Notice that the display groups the Level-1 and Level-2 LSPs separately.

Syntax: show isis database [detail | l1 | l2 | level1 | level2]

The **detail** parameter displays detailed information about the LSPs. See "Displaying Detailed Information" on page 6-35.

The **l1** and **level1** parameters display the Level-1 LSPs only. You can use either parameter. They do the same thing.

The **l2** and **level2** parameters display the Level-2 LSPs only. You can use either parameter. They do the same thing.

The **show isis database** summary display shows the following information.

Table 6.8: IS-IS Summary LSP Database Information

This Field...	Displays...
LSPID	The LSP ID, which consists of the source ID (6 bytes), the pseudonode (1 byte), and LSPID (1 byte). **Note**: If the address has an asterisk (*) at the end, this indicates that the LSP is locally originated.
LSP Seq Num	The sequence number of the LSP.

Table 6.8: IS-IS Summary LSP Database Information (Continued)

This Field...	Displays...
LSP Checksum	The checksum calculated by the device that sent the LSP and used by the Layer 3 Switch to verify that the LSP was not corrupted during transmission over the network.
LSP Holdtime	The maximum number of seconds during which the LSP will remain valid. **Note**: The IS that originates the LSP starts the timer for the LSP. As a result, LSPs do not all have the same amount of time remaining when they enter the Layer 3 Switch's LSP database.
ATT	A 4-bit value extracted from bits 4 – 7 in the Attach field of the LSP.
P	The value in the Partition option field of the LSP. The field can have one of the following values: • 0 – The IS that sent the LSP does not support partition repair. • 1 – The IS that sent the LSP supports partition repair.
OL	The value in the LSP database overload field of the LSP. The field can have one of the following values: • 0 – The overload bit is off. • 1 – The overload bit is on, indicating that the IS that sent the LSP is overloaded and should not be used as a Level-2 router.

Displaying Detailed Information

To display detailed information about the LSPs in the Layer 3 Switch's LSP databases, use either of the following methods.

USING THE CLI

To display detailed information for all the LSPs in the Layer 3 Switch's LSP databases, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis database detail

IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
00e0.5200.0001.00-00* 0x00000009   0x027b        1092              0/0/1
   Area Address:  20.8101
   NLPID:  cc
   IP address:  128.1.1.2
   Metric:  10    IP-Extended 128.1.1.0/24  UP bit: 0
   Metric:  10    IS 00e0.52b5.7800.01
   Metric:  10    IS-Extended 00e0.52b5.7800.01

LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
00e0.52b5.7800.00-00  0x00000007   0x8631        1024              0/0/0
   Area Address:  20.8101
   NLPID:  cc
   IP address:  128.1.1.1
   Metric:  10    IP-Internal 2.2.5.0          255.255.255.0
   Metric:  10    IP-Internal 2.2.6.0          255.255.255.0
   Metric:  10    IP-Internal 2.2.7.0          255.255.255.0
   Metric:  10    IP-Internal 2.2.8.0          255.255.255.0
   Metric:  10    IP-Internal 40.1.3.0         255.255.255.0
   Metric:  10    IP-Internal 128.1.1.0        255.255.255.0
   Metric:  10    IS 00e0.52b5.7800.01
   Metric:  10    IS 00e0.52b5.7800.02
```

Syntax: show isis database detail [l1 | l2 | level1 | level2]

The **detail** parameter displays detailed information about the LSPs. If you leave this parameter out, only summary information is displayed.

The **l1** and **level1** parameters display the Level-1 LSPs only. You can use either parameter. They do the same thing.

The **l2** and **level2** parameters display the Level-2 LSPs only. You can use either parameter. They do the same thing.

To display details about Level-1 or Level-2 LSPs only, use a combination of display options, as in the following example:

```
NetIron(config-isis-router)# show isis database level2 detail
```

This command displays detailed information for the Level-2 LSPs only.

The **show isis database detail** display shows the following information.

Table 6.9: IS-IS Detailed LSP Database Information

This Field...	Displays...
LSPID	See the description of the summary display.
LSP Seq Num	See the description of the summary display.
LSP Checksum	See the description of the summary display.

Table 6.9: IS-IS Detailed LSP Database Information (Continued)

This Field...	Displays...
LSP Holdtime	See the description of the summary display.
ATT/P/OL	See the description of the summary display.
Area Address	The address of the area.
NLPID	The Network Layer Protocol Identifier (NLPID), which specifies the protocol the IS that sent the LSP is using. Usually, this value is “cc” but can also be “iso”.
IP address	The IP address of the interface that sent the LSP. The Layer 3 Switch can use this address as the next hop in routes to the addresses listed in the rows below.
Destination addresses	The rows of information below the IP address row are the destinations advertised by the LSP. The Layer 3 Switch can reach these destinations by using the IP address listed above as the next hop. Each destination entry contains the following information: • Metric – The value of the default metric, which is the IS-IS cost of using the IP address above as the next hop to reach this destination. • Device type – The device type at the destination. The type can be one of the following: • End System – The device is an ES. • IP-Internal – The device is an ES within the current area. The IP address and sub-net mask are listed. • IS – The device is another IS. The NET (NSAP address) is listed. • IP-Extended – Same as IP-Internal, except the device uses the extended TLV fields described in draft-ietf-isis-traffic-02.txt to carry the information. • IS-Extended – Same as IS, except the device uses the extended TLV fields described in draft-ietf-isis-traffic-02.txt to carry the information.

Displaying Traffic Statistics

The Layer 3 Switch maintains statistics for common IS-IS PDU types. To display the statistics, use either of the following methods.

USING THE CLI

To display IS-IS PDU statistics, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis traffic
                             Message Received     Message Sent
  Level-1 Hellos             1029                 115
  Level-2 Hellos             1027                 112
  PTP Hellos                 0                    0
  Level-1 LSP                6                    3
  Level-2 LSP                6                    3
  Level-1 CSNP               0                    0
  Level-2 CSNP               0                    0
  Level-1 PSNP               107                  0
  Level-2 PSNP               107                  0
```

Syntax: show isis traffic

This display shows the following information.

Table 6.10: IS-IS Traffic Statistics

This Field...	Displays...
Level-1 Hellos	The number of Level-1 hello PDUs sent and received by the Layer 3 Switch.
Level-2 Hellos	The number of Level-2 hello PDUs sent and received by the Layer 3 Switch.
PTP Hellos	The number of point-to-point hello PDUs sent and received by the Layer 3 Switch.
Level-1 LSP	The number of Level-1 link-state PDUs sent and received by the Layer 3 Switch.
Level-2 LSP	The number of Level-2 link-state PDUs sent and received by the Layer 3 Switch.
Level-1 CSNP	The number of Level-1 Complete Sequence Number PDUs (CSNPs) sent and received by the Layer 3 Switch.
Level-2 CSNP	The number of Level-2 CSNPs sent and received by the Layer 3 Switch.
Level-1 PSNP	The number of Level-1 Partial Sequence Number PDUs (PSNPs) sent and received by the Layer 3 Switch.
Level-2 PSNP	The number of Level-2 PSNPs sent and received by the Layer 3 Switch.

Displaying Error Statistics

Use either of the following methods to display statistics for IS-IS errors.

USING THE CLI

To display IS-IS error statistics, enter the following command at any level of the CLI:

```
NetIron(config-isis-router)# show isis counts
 Area Mismatch: 0
 Max Area Mismatch: 0
 System ID Length Mismatch: 0
 Authentication Fail: 0
 Corrupted LSP: 0
 LSP Sequence Number Skipped: 0
 LSP Max Sequence Number Exceeded: 0
 Level-1 Database Overload: 0
 Level-2 Database Overload: 0
 Our LSP Purged: 0
```

Syntax: show isis counts

This display shows the following information.

Table 6.11: IS-IS Error Statistics

This Field...	Displays...
Area Mismatch	The number of times the Layer 3 Switch interface was unable to create a Level-1 adjacency with a neighbor because the Layer 3 Switch interface and the neighbor did not have any areas in common.
Max Area Mismatch	The number of times the Layer 3 Switch received a PDU whose value for maximum number of area addresses did not match the Layer 3 Switch’s value for maximum number of area addresses.
System ID Length Mismatch	The number of times the Layer 3 Switch received a PDU whose ID field was a different length than the ID field length configured on the Layer 3 Switch.
Authentication Fail	The Layer 3 Switch is configured to authenticate IS-IS packets in the packet’s domain or area, but the packet did not contain the correct password.
Corrupted LSP	The number of times the Layer 3 Switch detected a corrupted LSP in the device’s memory.
LSP Sequence Number Skipped	The number of times the Layer 3 Switch received an LSP with a sequence number that was more than 1 higher than the sequence number of the previous LSP received from the same neighbor.
LSP Max Sequence Number Exceeded	The number of times the Layer 3 Switch attempted to set an LSP sequence number to a value higher than the highest number in the CSNP sent by the Designated IS.

Table 6.11: IS-IS Error Statistics (Continued)

This Field...	Displays...
Level-1 Database Overload	The number of times the Level-1 state on the Layer 3 Switch changed from Waiting to On or from On to Waiting. • Waiting to On – This change can occur when the Layer 3 Switch recovers from a previous Level-1 LSP database overload and is again ready to receive new LSPs. • On to Waiting – This change can occur when the Layer 3 Switch's Level-1 LSP database is full and the Layer 3 Switch receives an additional LSP, for which there is no room.
Level-2 Database Overload	The number of times the Level-2 state on the Layer 3 Switch changed from Waiting to On or from On to Waiting. • The change from Waiting to On can occur when the Layer 3 Switch recovers from a previous Level-2 LSP database overload and is again ready to receive new LSPs. • The change from On to Waiting can occur when the Layer 3 Switch's Level-2 LSP database is full and the Layer 3 Switch receives an additional LSP, for which there is no room.
Our LSP Purged	The number of times the Layer 3 Switch received an LSP that was originated by the Layer 3 Switch itself and had age zero (aged out).

Clearing IS-IS Information

To clear the IS-IS information that the Layer 3 Switch has accumulated since the last time you cleared information or reloaded the software, use either of the following methods.

USING THE CLI

To clear IS-IS information, enter a command such as the following at any level of the CLI except the User EXEC level:

```
NetIron# clear isis all
```

This command clears all the following:

- Neighbors (closes the Layer 3 Switch's adjacencies with its IS-IS neighbors)
- Routes
- PDU statistics
- Error statistics

Syntax: clear isis all | counts | neighbor | route | traffic

The **all** parameter clears all the IS-IS information. Using this option is equivalent to entering separate commands with each of the other options.

The **counts** parameter clears the error statistics.

The **neighbor** parameter closes the Layer 3 Switch's adjacencies with its IS-IS neighbors and clears the neighbor statistics.

The **route** parameter clears the IS-IS route table.

The **traffic** parameter clears the PDU statistics.

NOTE: The **traffic** option also clears the values displayed in the **show isis interface** command’s Control Messages Sent and Control Messages Received fields.