

## Contents

|                                         |    |
|-----------------------------------------|----|
| Executive Summary .....                 | 3  |
| Demonstration Network Overview .....    | 4  |
| Routing Information Databases .....     | 7  |
| Using traffic-engineering bgp .....     | 11 |
| Using IGP Shortcuts .....               | 16 |
| Using traffic-engineering bgp-igp ..... | 22 |
| Installing Single Prefixes .....        | 28 |
| Appendix A .....                        | 31 |
| RTR1 .....                              | 31 |
| RTR2 .....                              | 32 |
| RTR3 .....                              | 33 |
| RTR4 .....                              | 34 |
| RTR5 .....                              | 35 |
| RTR6 .....                              | 35 |
| RTR7 .....                              | 36 |
| Acronyms .....                          | 37 |

## List of Figures

|                                                                                                        |    |
|--------------------------------------------------------------------------------------------------------|----|
| Figure 1: Demonstration Network Physical and Logical Topology .....                                    | 4  |
| Figure 2: Demonstration Network BGP Topology .....                                                     | 5  |
| Figure 3: Demonstration Network LSP Topology .....                                                     | 6  |
| Figure 4: By Default, BGP Uses Both inet.0 and inet.3 .....                                            | 11 |
| Figure 5: IGP Shortcuts Allow the IGP to Install Prefixes in inet.3 .....                              | 18 |
| Figure 6: The traffic-engineering bgp-igp Command Causes RSVP to<br>Install Prefixes into inet.0 ..... | 22 |

## List of Tables

|                                                                             |   |
|-----------------------------------------------------------------------------|---|
| Table 1: Routing Tables for MPLS Traffic Engineering Route Resolution ..... | 7 |
|-----------------------------------------------------------------------------|---|

## Executive Summary

Multiprotocol Label Switching (MPLS) traffic engineering maps certain data flows to established label switched paths (LSPs) rather than to data links calculated by the IGP to be part of the *best* (shortest) path. Fundamental to this function is the determination of what traffic is to be mapped to an LSP.

Traffic is mapped to an LSP at the tunnel's ingress label switching router (LSR) by designating the egress LSR as the next-hop router for certain destination prefixes. It is important to understand that the LSP does not constitute an entire route to a destination. Rather, the LSP is a next-hop segment of the route. Therefore, packets can only be mapped to an LSP if the egress LSR is considered to be a feasible next-hop candidate during the route resolution process. This paper explains the various options available with Juniper Networks JUNOS<sup>™</sup> Internet software for including LSPs in the route resolution process.


This paper assumes that you understand the basics of MPLS and the JUNOS command-line interface.

## Demonstration Network Overview

The same network topology is used for all of the route resolution examples. This section provides reference diagrams necessary for the understanding of the examples that are presented in subsequent sections.


Figure 1 shows the routers, router names, physical connections, and IP addresses used in the network. There are two autonomous systems and seven routers: RTR1 through RTR7. The loopback interface addresses serve as the OSPF and BGP router IDs, and as the endpoints for IBGP and LSP connections, where appropriate.

**Figure 1: Demonstration Network Physical and Logical Topology**




The IGP used in AS 65501 is OSPF, and all internal interfaces are in area 0. The external interface at RTR4 linking to AS 65520 does not run OSPF. Figure 2 shows the logical BGP topology. A full IBGP mesh is configured within AS 65501, linking all routers except RTR6. No BGP runs on that router. An EBGP connection links RTR4 in AS 65501 with RTR7 in AS 65520. The endpoints of the EBGP connection are the external physical interfaces of the two routers.

**Figure 2: Demonstration Network BGP Topology**



A single LSP, named `test_path`, is configured from RTR1 to RTR4, as depicted in Figure 3. The LSP Explicit Route Object (ERO) is specified to use a strict hop through RTR2, so that the LSP takes a different path from the OSPF shortest path of RTR1-RTR5-RTR4. The LSP is signaled using RSVP, but no CSPF is running.

**Figure 3: Demonstration Network LSP Topology**



Refer to Appendix A for the base routing options, policy, and protocol configurations for all seven routers.

## Routing Information Databases

JUNOS software maintains several routing tables from which best-path information is extracted and added to the forwarding table. Both the maintenance and use of the routing tables is protocol dependent. The routing tables relevant to MPLS traffic engineering route resolution are `inet.0`, `inet.3`, and `mpls.0`.

| Table 1: Routing Tables for MPLS Traffic Engineering Route Resolution |                       |                                                                                                                                           |
|-----------------------------------------------------------------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Routing Table                                                         | Type                  | Description                                                                                                                               |
| <code>inet.0</code>                                                   | Unicast routing table | All unicast protocols enter their route information into this table.                                                                      |
| <code>inet.3</code>                                                   | MPLS routing table    | RSVP enters destinations reachable via LSPs into this table; these destination prefixes are always the IP addresses of LSP egress points. |
| <code>mpls.0</code>                                                   | MPLS switching table  | MPLS enters label bindings into this table, and uses the information to switch incoming MPLS-labeled packets to outgoing interfaces.      |

With the default configurations listed in Appendix A running, RTR1's three routing tables are as follows.

```
jeff@RTR1> show route
inet.0: 20 destinations, 20 routes (18 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.1.0/24      *[Direct/0] 17:34:14
                   > via fxp0.0
192.168.1.1/32     *[Local/0] 17:34:14
                   Local
192.168.2.0/24     *[Direct/0] 17:34:14
                   > via fxp1.0
192.168.2.1/32     *[Local/0] 17:34:14
                   Local
192.168.3.0/24     *[OSPF/10] 17:33:27, metric 2
                   > to 192.168.1.2 via fxp0.0
192.168.4.0/24     *[OSPF/10] 17:33:22, metric 2
                   > to 192.168.2.2 via fxp1.0
192.168.5.0/24     *[OSPF/10] 17:33:22, metric 3
                   to 192.168.1.2 via fxp0.0
                   > to 192.168.2.2 via fxp1.0
192.168.6.0/24     *[OSPF/10] 17:33:22, metric 3
                   > to 192.168.2.2 via fxp1.0
192.168.7.0/24     *[OSPF/10] 17:33:22, metric 13
                   > to 192.168.2.2 via fxp1.0
192.168.8.0/24     *[OSPF/10] 17:33:22, metric 13
                   > to 192.168.2.2 via fxp1.0
192.168.9.0/24     *[OSPF/10] 17:33:22, metric 12
                   > to 192.168.2.2 via fxp1.0
192.168.254.1/32   *[Direct/0] 17:34:14
                   > via lo0.0
192.168.254.2/32   *[OSPF/10] 17:33:27, metric 1
                   > to 192.168.1.2 via fxp0.0
192.168.254.3/32   *[OSPF/10] 17:33:27, metric 2
                   > to 192.168.1.2 via fxp0.0
192.168.254.4/32   *[OSPF/10] 17:33:22, metric 2
                   > to 192.168.2.2 via fxp1.0
192.168.254.5/32   *[OSPF/10] 17:33:22, metric 1
                   > to 192.168.2.2 via fxp1.0
192.168.254.6/32   *[OSPF/10] 17:33:22, metric 3
                   > to 192.168.2.2 via fxp1.0
224.0.0.5/32       *[OSPF/10] 17:34:15, metric 1

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.254.4/32   *[RSVP/7] 17:33:22, metric 2, metric2 0
                   > to 192.168.1.2 via fxp0.0, label-switched-path test_path

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 17:34:14, metric 1
                   Receive
1                  *[MPLS/0] 17:34:14, metric 1
                   Receive
```

The first table displayed is `inet.0`, which contains 18 active prefixes. The table here shows that with the exception of local and directly connected subnets, all prefixes were entered by OSPF.

The second table displayed is `inet.3`, containing a single entry. The destination prefix is `192.168.254.4/32`, which is the egress of LSP `test_path`, and is reachable via that LSP.

The last table is `mpls.0`. This table contains entries for labels 0 and 1, which are entered automatically by MPLS when the protocol is enabled. These labels, when received, signify an LSP endpoint and a router alert, respectively. No other labels are entered into `mpls.0` because RTR1 is not a transit router for any LSP. Contrast RTR1's `mpls.0` table with that of RTR2.

```
jeff@RTR2> show route table mpls.0
```

```
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
0          *[MPLS/0] 20:12:23, metric 1
            Receive
1          *[MPLS/0] 20:12:23, metric 1
            Receive
100004     *[RSVP/7] 18:07:14, metric 1
            > to 192.168.3.2 via fxp1.0, label-switched-path test_path
```

This table shows that incoming MPLS packets with a label of 100004 are switched to interface fxp1, following LSP test\_path. You can view the full path of the LSP, including the label bindings, by displaying the RSVP sessions on each of the four routers in the path.

```
jeff@RTR1> show rsvp session brief
Ingress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
192.168.254.4 192.168.254.1 Up    0   1 FF      -    100004 test_path
Total 1 displayed, Up 1, Down 0
Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
jeff@RTR2> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
192.168.254.4 192.168.254.1 Up    1   1 FF  100004  100004 test_path
Total 1 displayed, Up 1, Down 0
```

```
jeff@RTR3> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
192.168.254.4 192.168.254.1 Up    1   1 FF  100004      3 test_path
Total 1 displayed, Up 1, Down 0
```


```
jeff@RTR4> show rsvp session brief
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
192.168.254.4 192.168.254.1 Up    0   1 FF      3      - test_path
Total 1 displayed, Up 1, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

On all four routers, the ingress and egress IP addresses of the LSP are shown. The path is shown as an ingress path at RTR1, and packets forwarded on the LSP are assigned a label of 100004. At RTR2 the LSP is transit, and packets arriving with a label of 100004 are given an outgoing label of 100004. Although these two labels are numerically the same, label swapping still is taking place at RTR2; the labels have significance only between neighboring LSRs. RTR3 swaps incoming label 100004 for outgoing label 3. RTR4, the egress, pops label 3 and routes the received packet by standard IP longest-match route lookup.

## Using traffic-engineering bgp

When MPLS is enabled, the command `traffic-engineering bgp` is enabled by default. Since it is a default, it is implicit and does not appear in the JUNOS MPLS configuration stanzas. The `traffic-engineering bgp` command causes BGP to examine both the `inet.0` and the `inet.3` tables when resolving a route, as shown in Figure 4. All other unicast routing protocols consult only `inet.0`. The significance of this behavior is not with the resolution of destination addresses, but with the resolution of next-hop addresses, as the example in this section demonstrates.

**Figure 4: By Default, BGP Uses Both `inet.0` and `inet.3`**



以上内容仅为本文档的试下载部分，为可阅读页数的一半内容。如要下载或阅读全文，请访问：<https://d.book118.com/867053134014006036>