MPLS Traffic Engineering--Scalability Enhancements

The MPLS Traffic Engineering--Scalability Enhancement feature improves scalability performance for large numbers of traffic engineering tunnels.

These improvements allow an increase in the number of traffic engineering (TE) tunnels a router can support when the router is configured as a tunnel headend. Additionally, when the router is configured as a tunnel midpoint, the enhancements reduce the time required to establish large numbers of TE tunnels.

This feature module contains information about and instructions on how to configure the Multiprotocol Label Switching (MPLS) traffic engineering scalability enhancements.

- Finding Feature Information, page 1
- Prerequisites for MPLS Traffic Engineering--Scalability Enhancements, page 2
- Restrictions for MPLS Traffic Engineering--Scalability Enhancements, page 2
- Information About MPLS Traffic Engineering--Scalability Enhancements, page 2
- How to Configure MPLS Traffic Engineering--Scalability Enhancements, page 4
- Configuration Examples for MPLS Traffic Engineering--Scalability Enhancements, page 11
- Additional References, page 12
- Feature Information for MPLS Traffic Engineering Scalability Enhancements, page 14
- Glossary, page 15

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for MPLS Traffic Engineering--Scalability Enhancements

Your network must support the following Cisco IOS XE features before you enable MPLS traffic engineering:

- MPLS
- Cisco Express Forwarding
- Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF)

Restrictions for MPLS Traffic Engineering--Scalability Enhancements

The number of tunnels that a particular platform can support can vary depending on:

- The types of interfaces that the tunnels traverse
- The manner in which the Resource Reservation Protocol (RSVP) message pacing feature is configured
- MPLS traffic engineering supports only a single IGP process/instance. Multiple IGP processes/instances are not supported and MPLS traffic engineering should not be configured in more than one IGP process/instance.

Information About MPLS Traffic Engineering--Scalability Enhancements

Scalability Enhancements for Traffic Engineering Tunnels

Scalability performance is improved for large numbers of traffic engineering tunnels, and includes the following enhancements:

- Increase the number of traffic engineering tunnels a router can support when configured as a tunnel headend and when configured as a tunnel midpoint
- Reduce the time required to establish large numbers of traffic engineering tunnels

RSVP Rate Limiting

A burst of RSVP traffic engineering signaling messages can overflow the input queue of a receiving router, causing some messages to be dropped. Dropped messages cause a substantial delay in completing label switched path (LSP) signaling.

This MPLS Traffic Engineering--Scalability Enhancements feature provides an enhancement mechanism that controls the transmission rate for RSVP messages and reduces the likelihood of input drops on the receiving router. The default transmission rate is 200 RSVP messages per second to a given neighbor. The rate is configurable.

Improved Recovery Response for Signaling and Management of MPLS Traffic Engineering Tunnels

The MPLS Traffic Engineering--Scalability Enhancements feature improves the recovery response for signaling and management of MPLS TE tunnels. LSP recovery responsiveness is improved when a link used by an LSP fails:

- When the upstream end of a failed link detects the failure, the software generates an RSVP No Route path error message. This enables the LSP headend to detect the link failure and initiate recovery, even when the Interior Gateway Protocol (IGP) update announcing the link failure is delayed.
- The LSP headend marks the link in question so that subsequent constraint-based shortest path first (SPF) calculations ignore the link until either a new IGP update arrives or a configurable timeout occurs. This ensures that resignaling to restore the LSP avoids the failed link.

IS-IS and MPLS Traffic Engineering Topology Database Interactions

The MPLS Traffic Engineering--Scalability Enhancements feature reduces the interval between when the IS-IS protocol receives an IGP update and when it delivers the update to the MPLS traffic engineering topology database.

Before the MPLS Traffic Engineering--Scalability Enhancements feature was introduced, when IS-IS received a new LSP that contained traffic engineering type, length, value (TLV) objects, a delay of several seconds could occur before IS-IS passed the traffic engineering TLVs to the traffic engineering database. The purpose of the delay was to provide better scalability during periods of network instability and to give the router an opportunity to receive more fragments of the LSP before passing the information to the traffic engineering database. However, this delay increased the convergence time for the traffic engineering database.

With the MPLS Traffic Engineering--Scalability Enhancements feature, IS-IS extracts traffic engineering TLVs from received LSPs and passes them to the traffic engineering database immediately. The exception to this occurs when there are large numbers of LSPs to process and it is important to limit CPU consumption, such as during periods of network instability. The parameters that control IS-IS delivery of traffic engineering TLVs to the traffic engineering topology database are configurable.

Note MPLS traffic engineering supports only a single IGP process/instance. Multiple IGP processes/instances are not supported and MPLS traffic engineering should not be configured in more than one IGP process/instance.

Improved Counter Capabilities for MPLS TE Tunnels Events and RSVP Signaling

With the MPLS Traffic Engineering--Scalability Enhancements feature, diagnostic and troubleshooting capabilities for MPLS traffic engineering tunnels and RSVP are improved:

- Counters record tunnel headend error events such as no route (link down), preemption, and insufficient bandwidth on a per-tunnel basis.
- Counters record RSVP messages. The counters are per-interface and record the number of RSVP messages of each type sent and received on the interface.

Benefits of MPLS Traffic Engineering--Scalability Enhancements

The MPLS Traffic Engineering--Scalability Enhancements feature provides the following benefits:

- Increased scalability--Up to 600 MPLS traffic engineering tunnel headends are supported. Up to 10,000 traffic engineering tunnel midpoints are supported, with up to 5000 midpoints per interface.
- Faster recovery after failure conditions--Message pacing provides a mechanism to throttle RSVP control messages so that they are less likely to be dropped. This results in a faster recovery from failure conditions when many MPLS traffic engineering tunnels are being set up.
- Improved reroute time--When a traffic engineering tunnel is down, the headend router needs to be notified so that it can signal for a new LSP for the tunnel along an alternate path. The headend router does not have to wait for an IGP update to signal for a new LSP for the tunnel along an alternate path.
- Improved tunnel setup time--Fewer control messages and tunnel setup messages are dropped. This reduces the average time required to set up tunnels.

How to Configure MPLS Traffic Engineering--Scalability Enhancements

Enabling RSVP Rate Limiting for MPLS Traffic Engineering Scalability Enhancements

Perform the following task to enable RSVP rate limiting for MPLS traffic engineering scalability enhancements. RSVP rate limiting maintains, on an outgoing interface basis, a count of messages that were dropped because the output queue for the interface used for rate limiting was full.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip rsvp signalling rate-limit** [**burst** *number*] [**limit** *number*] [**maxsize** *bytes*] [**period** *ms*]
4. **end**
5. **show ip rsvp neighbor**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**ip rsvp signalling rate-limit** [**burst** *number*] [**limit** *number*] [**maxsize** *bytes*] [**period** *ms*] **Example:** `Router(config)# ip rsvp signalling rate-limit burst 5 maxsize 3 period 2`	Controls the transmission rate for RSVP messages sent to a neighboring router during a specified amount of time. • The **burst** *number*keyword and argument pair indicates the maximum number of RSVP messages sent to a neighboring router during each interval. The range is from 1 to 5000. The default is 8. • The **limit** *number* keyword and argument pair indicates the maximum number of messages to send per queue interval when the number of messages sent is less than the number of messages to be sent normally. The range is 1 to 5000. The default is 37. • The **maxsize** *bytes* keyword and argument pair indicates the maximum size of the message queue, in bytes. The range is 1 to 5000. The default is 2000. • The **period** *ms* keyword and argument pair indicates the length of the interval (time frame) in milliseconds (ms). The range is 10 to 5000. The default is 20.
Step 4	**end** **Example:** `Router(config)# end`	Exits to privileged EXEC mode.
Step 5	**show ip rsvp neighbor** **Example:** `Router# show ip rsvp neighbor`	Displays current RSVP neighbors. Use this command to verify that RSVP message pacing is enabled.

Managing Link Failure Timeouts for MPLS Traffic Engineering Tunnels

Perform this task to manage link failure timeouts for MPLS traffic engineering tunnels.

This allows the configuration of a timeout during which the router ignores a link in its path calculation to avoid paths that contain a failed link and are likely to fail when signaled.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **mpls traffic-eng topology holddown sigerr** *seconds*
4. **end**
5. **show mpls traffic-eng topology** [**brief**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**mpls traffic-eng topology holddown sigerr** *seconds* **Example:** `Router(config)# mpls traffic-eng topology holddown sigerr 15`	Specifies the amount of time that a router ignores a link in its traffic engineering topology database in tunnel path Constrained Shortest Path First (CSPF) computations following a traffic engineering tunnel error on the link. • The *seconds* argument specifies the length of time (in seconds) a router should ignore a link during tunnel path calculations following a traffic engineering tunnel error on the link. The range is 0 to 300. The default is 10.
Step 4	**end** **Example:** `Router(config)# end`	Exits to privileged EXEC mode.

	Command or Action	Purpose
Step 5	**show mpls traffic-eng topology** [**brief**] **Example:** `Router# show mpls traffic-eng topology brief`	Displays the MPLS traffic engineering global topology as currently known at this node. • The **brief** keyword provides a less detailed version of the topology.

Controlling IS-IS Communication with the MPLS Traffic Engineering Topology Database

Perform the following task to control IS-IS and MPLS traffic engineering topology database interactions. This reduces the interval time between when the IS-IS protocol receives an IGP update and when IS-IS delivers the update to the MPLS traffic engineering topology database, which reduces convergence time for the database.

Note MPLS traffic engineering supports only a single IGP process/instance. Multiple IGP processes/instances are not supported and MPLS traffic engineering should not be configured in more than one IGP process/instance.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router isis** [*area-tag*]
4. **mpls traffic-eng scanner** [**interval** *seconds*] [**max-flash** *LSPs*]
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**router isis** [*area-tag*] **Example:** `Router(config)# router isis`	Enables the IS-IS routing protocol and specifies an IS-IS process. • The *area-tag* argument is a meaningful name for a routing process. If it is not specified, a null tag is assumed and the process is referenced with a null tag. This name must be unique among all IP or Connectionless Network Service (CLNS) router processes for a given router. **Note** This argument is Required for multiarea IS-IS configuration and optional for conventional IS-IS configuration.
Step 4	**mpls traffic-eng scanner** [**interval** *seconds*] [**max-flash** *LSPs*] **Example:** `Router(config-router)# mpls traffic-eng scanner interval 5 max-flash 100`	Specifies how often IS-IS extracts traffic engineering TLVs from flagged LSPs and passes them to the traffic engineering topology database, and specifies the maximum number of LSPs that the router can process immediately. • The **interval** *seconds* keyword and argument specify the frequency, in seconds, at which IS-IS sends traffic engineering TLVs into the traffic engineering database. The range is 1 to 60. The default is 5. • The **max-flash** *LSPs* keyword and argument specify the maximum number of LSPs that the router can process immediately without incurring a delay. The range is 0 to 200. The default is 15.
Step 5	**end** **Example:** `Router(config-router)# end`	Exits to privileged EXEC mode.

Monitoring and Maintaining MPLS TE Scalability Enhancements

SUMMARY STEPS

1. **enable**
2. **show ip rsvp neighbor** [**detail**]
3. **show ip rsvp counters** [**summary**]
4. **clear ip rsvp counters**
5. **clear ip rsvp signalling rate-limit**
6. **show mpls traffic-eng tunnels statistics**
7. **clear mpls traffic-eng tunnels counters**
8. **show mpls traffic-eng topology** [**brief**]
9. **exit**

DETAILED STEPS

Step 1 **enable**

Use this command to enable privileged EXEC mode. Enter your password if prompted. For example:

Example:

```
Router> enable
Router#
```

Step 2 **show ip rsvp neighbor** [**detail**]

Use this command to verify that RSVP message pacing is turned on. For example:

Example:

```
Router# show ip rsvp neighbor detail
Neighbor:10.0.0.1
    Encapsulation:RSVP
    Rate-Limiting:
      Dropped messages:0
Refresh Reduction:
      Remote epoch:0x1BFEA5
      Out of order messages:0
      Retransmitted messages:0
      Highest rcvd message id:1059
      Last rcvd message:00:00:04
Neighbor:10.0.0.2
    Encapsulation:RSVP
    Rate-Limiting:
      Dropped messages:0
    Refresh Reduction:
      Remote epoch:0xB26B1
      Out of order messages:0
      Retransmitted messages:0
      Highest rcvd message id:945
      Last rcvd message:00:00:05
```

Step 3 **show ip rsvp counters** [**summary**]

Use this command to display the counts of RSVP messages that were sent and received. For example:

Example:

```
Router# show ip rsvp counters summary
All Interfaces          Recv      Xmit                        Recv      Xmit
    Path                  110        15    Resv                    50        28
    PathError               0         0    ResvError               0         0
    PathTear                0         0    ResvTear                0         0
    ResvConf                0         0    RTearConf               0         0
    Ack                     0         0    Srefresh                0         0
    Hello                5555      5554    IntegrityChalle         0         0
    IntegrityRespon         0         0    DSBM_WILLING            0         0
    I_AM_DSBM               0         0
    Unknown                 0         0    Errors                  0         0
Recv Msg Queues                Current       Max
    RSVP                             0         2
    Hello (per-I/F)                  0         1
    Awaiting Authentication          0         0
```

Step 4 **clear ip rsvp counters**

Use this command to clear (set to zero) all IP RSVP counters that are being maintained. For example:

Example:

```
Router# clear ip rsvp counters
Clear rsvp counters [confirm]
```

Step 5 **clear ip rsvp signalling rate-limit**
Use this command to clear (set to zero) counts of the messages that message pacing was forced to drop because the output queue for the interface used for message pacing was full. For example:

Example:

```
Router# clear ip rsvp signalling rate-limit
```

Step 6 **show mpls traffic-eng tunnels statistics**
Use this command to display event counters for one or more MPLS traffic engineering tunnels. For example:

Example:

```
Router# show mpls traffic-eng tunnels statistics
Tunnel1001 (Destination 10.8.8.8; Name Router_t1001)
  Management statistics:
    Path: 25 no path, 1 path no longer valid, 0 missing ip exp path
 5 path changes
    State: 3 transitions, 0 admin down, 1 oper down
  Signalling statistics:
    Opens: 2 succeeded, 0 timed out, 0 bad path spec
 0 other aborts
    Errors: 0 no b/w, 0 no route, 0 admin
 0 bad exp route, 0 rec route loop, 0 other

...
```

Example:

```
Tunnel7050 (Destination 10.8.8.8; Name Router_t7050)
  Management statistics:
    Path:  19 no path, 1 path no longer valid, 0 missing ip exp path
 3 path changes
    State:  3 transitions, 0 admin down, 1 oper down
  Signalling statistics:
    Opens:  2 succeeded, 0 timed out, 0 bad path spec
 0 other aborts
    Errors: 0 no b/w, 0 no route, 0 admin
 0 bad exp route, 0 rec route loop, 0 other
```

Step 7 **clear mpls traffic-eng tunnels counters**
Use this command to clear counters for all MPLS traffic engineering tunnels. For example:

Example:

```
Router# clear mpls traffic-eng tunnels counters
Clear traffic engineering tunnel counters [confirm]
```

Step 8 **show mpls traffic-eng topology** [**brief**]
Use this command to display the MPLS traffic engineering topology database. For example:

Example:

```
Router# show mpls traffic-eng topology brief
My_System_id:0000.0000.0003.00 (isis  level-2)
Signalling error holddown:10 sec Global Link Generation 9
IGP Id:0000.0000.0003.00, MPLS TE Id:10.0.3.1 Router Node  (isis
level-2)
      link[0]:Point-to-Point, Nbr IGP Id:0000.0000.0004.00,
nbr_node_id:2, gen:9
          frag_id 0, Intf Address:10.0.0.33, Nbr Intf Address:10.0.0.34
          TE metric:10, IGP metric:10, attribute_flags:0x0
          SRLGs:1 2
```

Step 9 **exit**

Use this command to exit to user EXEC mode. For example:

Example:

```
Router# exit
Router>
```

Configuration Examples for MPLS Traffic Engineering--Scalability Enhancements

Example Enabling RSVP Rate Limiting for MPLS Traffic Engineering Scalability Enhancements

The following examples show how to enable RSVP rate limiting for MPLS traffic engineering scalability enhancements:

```
configure terminal
 ip rsvp signalling rate-limit
 end
```

The following is sample output that traffic engineering displays when RSVP rate limiting is enabled:

```
Router# show ip rsvp signalling rate-limit
Rate Limiting: enabled
  Burst: 10
  Limit: 37
  Maxsize: 5000
  Period (msec): 100
  Max rate (msgs/sec): 100
```

The following example shows how to configure a router to send a maximum of 5 RSVP traffic engineering signaling messages in 1 second to a neighbor. The size of the output queue is 35.

```
configure terminal
 ip rsvp signalling rate-limit
period 1 burst 5 maxsize 35
```

Example Managing Link Failure Timeouts for MPLS Traffic Engineering Tunnels

The following example shows how to manage link failure timeouts for MPLS traffic engineering tunnels:

```
configure terminal
 mpls traffic-eng topology holddown sigerr 15
 end
```

In this example, the link hold-down time for signaling errors is set to 15 seconds.

Example Controlling IS-IS Communication with the MPLS Traffic Engineering Topology Database

The following example shows how to control IS-IS communication with the MPLS traffic engineering topology database:

```
configure terminal
router isis
 mpls traffic-eng scanner interval 5 max-flash 50
 end
```

In this example, the router is enabled to process up to 50 IS-IS LSPs without any delay.

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
Quality of service	• *Cisco IOS Quality of Service Solutions Command Reference* • *Cisco IOS XE Quality of Service Solutions Configuration Guide,* Release 2
MPLS	• *Cisco IOS Multiprotocol Label Switching Command Reference* • *Cisco IOS XE Multiprotocol Label Switching Configuration Guide,* Release 2

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	--

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	--

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for MPLS Traffic Engineering Scalability Enhancements

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for MPLS Traffic Engineering Scalability Enhancements

Feature Name	Releases	Feature Information
MPLS Traffic Engineering: Scalability Enhancements	Cisco IOS XE Release 2.3	The MPLS Traffic Engineering--Scalability Enhancements feature improves scalability performance for large numbers of traffic engineering tunnels. These improvements allow an increase in the number of traffic engineering (TE) tunnels a router can support when the router is configured as a tunnel headend. Additionally, when the router is configured as a tunnel midpoint, the enhancements reduce the time required to establish large numbers of TE tunnels. This feature module contains information about and instructions on how to configure the Multiprotocol Label Switching (MPLS) traffic engineering scalability enhancements. The following commands were introduced or modified: **clear ip rsvp counters**, **clear ip rsvp signalling rate-limit**, **clear mpls traffic-eng tunnel counters**, **ip rsvp signalling rate-limit**, **mpls traffic-eng scanner**, **mpls traffic-eng topology holddown sigerr**, **show ip rsvp counters**, and **show mpls traffic-eng tunnels statistics**.

Glossary

bundled interface—Generic terms to represent port-channel, multilink, and VLAN interfaces.

Cisco express forwarding —A means for accelerating the forwarding of packets within a router, by storing route lookup information in several data structures instead of in a route cache.

CLNS —Connectionless Network Service. The Open Systems Interconnection (OSI) network layer service that does not require a circuit to be established before data is transmitted. CLNS routes messages to their destination independently of any other messages.

CSPF —Constrained Shortest Path First. A routing protocol that calculates the shortest path based on a set of constraints, such as a minimum bandwidth requirement, maximum number of nodes, or nodes to include or exclude.

enterprise network —A large and diverse network connecting most major points in a company or other organization.

FRR—Fast ReRoute.

headend —The endpoint of a broadband network. All stations send toward the headend; the headend then sends toward the destination stations.

IGP —Interior Gateway Protocol. An Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include Interior Gateway Routing protocol (IGRP), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP).

interface —A network connection.

IS-IS —Intermediate System to Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, where ISs (routers) exchange routing information based on a single metric, to determine the network topology.

LDN— Link Down Notification.

LSP —Label-Switched Path. A sequence of hops (R0...Rn) in which a packet travels from R0 to Rn through label switching mechanisms. A label-switched path can be chosen dynamically, based on normal routing mechanisms, or through configuration.

member links—Individual interfaces that are grouped into a bundled interface.

message-pacing —The former name of the rate limiting feature.

MPLS —Formerly known as tag switching, Multiprotocol Label Switching is a method for directing packets primarily through Layer 2 switching rather than Layer 3 routing. In MPLS, packets are assigned short fixed-length labels at the ingress to an MPLS cloud by using the concept of forwarding equivalence classes. Within the MPLS domain, the labels are used to make forwarding decisions mostly without recourse to the original packet headers.

OSPF —Open Shortest Path First. A link-state, hierarchical Interior Gateway Protocol (IGP) routing protocol. derived from the Intermediate System-Intermediate System (IS-IS) protocol. OSPF features are least-cost routing, multipath routing, and load balancing.

router —A network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information.

RSVP —Resource Reservation Protocol. A protocol that supports the reservation of resources across an IP network.

scalability —An indicator showing how quickly some measure of resource usage increases as a network gets larger.

TLV —type, length, value. TLV objects are used in data communication to provide optional information. The type field indicates the type of items in the value field. The length field indicates the length of the value field. The value field is the data portion of the packet.

topology —The physical arrangement of network nodes and media within an enterprise networking structure.

TE (traffic engineering) —Techniques and processes that cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods were used.

traffic engineering tunnel —A label-switched tunnel that is used for traffic engineering. Such a tunnel is set up through means other than normal Layer 3 routing; it is used to direct traffic over a path different from the one that Layer 3 routing would cause the tunnel to take.