



# Configuring MPLS Segment Routing OAM

This chapter describes the Multiprotocol Label Switching (MPLS) segment routing OAM functionality.

- [Overview of MPLS Segment Routing OAM, on page 1](#)
- [Segment Routing OAM Support for LSP Ping and Traceroute , on page 1](#)
- [Guidelines and Limitations for MPLS OAM, on page 2](#)
- [Examples for Using Ping and Traceroute CLI Commands, on page 3](#)

## Overview of MPLS Segment Routing OAM

BGP MPLS segment routing (SR) has been deployed on the Cisco Nexus 9000 Series switches. As MPLS segment routing (SR) is deployed, a few diagnostic tools are required to help resolve the misconfigurations or failures in the segment routing network. Beginning with Cisco NX-OS Release 7.0(3)I6(1), MPLS OAM is supported on the Cisco Nexus 9000 Series switches. In this first introduction, only Nil FEC is supported and none of the other FEC types are supported. The Nil FEC is the basic OAM FEC that is described in RFC-4379.

MPLS OAM provides two main functions for diagnostics purposes:

1. MPLS ping
2. MPLS traceroute

OAM draws the information from the FEC type to help diagnose the issues. The Nil FEC is not associated with a protocol like the other FEC types, and it is also not associated with a real FEC. For example, it is not associated with LDP etc. Logically, it only validates the data plane programming; it does not query the BGP or other routing protocols in the control plane unlike other FEC types.

To enable MPLS OAM on Cisco Nexus 9000 Series switches, use the **feature mpls oam** CLI command. Use the **no feature mpls oam** CLI command to disable MPLS OAM on Cisco Nexus 9000 Series switches.

## Segment Routing OAM Support for LSP Ping and Traceroute

The Nil-FEC LSP ping and traceroute operations are extensions of regular MPLS ping and traceroute. Nil-FEC LSP Ping/Traceroute functionality supports segment routing and MPLS Static. It also acts as an additional diagnostic tool for all other LSP types. This feature allows operators to provide the ability to freely test any label stack by allowing them to specify the following:

- Label stack
- Outgoing interface
- Nexthop address

In case of segment routing, each segment nodal label and adjacent label along the routing path is put into the label stack of an echo request message from the initiator Label Switch Router (LSR); MPLS data plane forwards this packet to the label stack target, and the label stack target sends the echo message back.

Use the **ping mpls nil-fec labels comma-separated-labels [output {interface tx-interface} [nexthop nexthop-ip-addr]]** CLI command to execute a ping. Use the **traceroute mpls nil-fec labels comma-separated-labels [output {interface tx-interface} [nexthop nexthop-ip-addr]]** CLI command to execute a traceroute.

## Guidelines and Limitations for MPLS OAM

See the following guidelines and limitations for configuring MPLS OAM Nil FEC:

- For notes on platform support see: [Platform Support for Label Switching Features](#).
- A maximum of four labels can be specified in the **ping mpls nil-fec** and **traceroute mpls nil-fec** commands. This value is enforced by querying the platform and currently Cisco Nexus 9000 Series switches limit the label stack to 5. It means that for a Nil FEC echo request, you can specify a maximum of four labels because internally an extra explicit-null is added.
- The next hop specified in the ping and traceroute commands must be a connected next hop on the originator and it should not be a recursive next hop.
- There is no support for tree trace.
- Nil FEC does not carry any information to identify the intended target. The packet may mis-forward at an incorrect node but the validation may return success if the packet ends up at a node after popping the non-null labels.
- Nil FEC operates on forwarding the information alone. It cannot detect the inconsistencies between the control plane and the forwarding plane by definition.
- Nil FEC ping and traceroute is not supported for deaggregator (per-VRF) labels. This includes the BGP EVPN-Layer three deaggregator labels.
- On Cisco Nexus 9000 Series switches that use Broadcom chipsets, there is no support to allow the software to send a query to determine which ECMP a packet takes. It means that for MPLS traceroutes that traverse one of these switches may display an error at the next hop if there is more than one ECMP as displayed in the following example:
 

```
D 2 6.0.0.2 MRU 1496 [Labels: 2003/explicit-null Exp: 0/0] 4 ms
```
- When you use OAM to test a BGP EPE LSP (for example, the last label in the ping or traceroute label stack is an EPE label), OAM only returns success if the final router has OAM enabled and MPLS is enabled on the incoming interface.

For example, if you have a setup as A---B---C, A and B are in the SR network, and B acts like a PE and C acts like a CE, B is configured with C as a BGP EPE peer (using egress-engineering on B), then C must have OAM and MPLS forwarding enabled on the incoming interface.

# Examples for Using Ping and Traceroute CLI Commands

## Using CLI to Execute a Ping

Use the **ping mpls nil-fec labels** *comma-separated-labels* [**output** {**interface** *tx-interface*} [**nexthop** *nexthop-ip-addr*]] CLI command to execute a ping.

For example, the following command sends an MPLS packet with the outermost two labels in the label stack being 2001 and 2000 out the interface Ethernet 1/1 with a nexthop IP address of 4.0.0.2:

```
switch# ping mpls nil-fec labels 2001,2000 output interface e1/1 nexthop 4.0.0.2
```

It is mandatory that the nexthop is a connected nexthop; it is not recursively resolved.

The above CLI format is a simplified version. The [**output** {**interface** *tx-interface*} [**nexthop** *nexthop-ip-addr*]] is mandatory to be present in the VSH server. For example:

```
switch# ping mpls nil-fec labels 1,2 ?
output Output options
```

```
switch# ping mpls nil-fec labels 1,2
^
% Invalid command at '^' marker.
```

## Using CLI to Execute a Traceroute

Use the following CLI command to execute a traceroute:

```
traceroute mpls nil-fec labels <comma-separated-labels> output interface <tx-interface>
nexthop <nexthop-ip-addr>
```

## Displaying Show Statistics

Use the following command to display the statistics about the echo requests sent by the local MPLS OAM service:

```
show mpls oam echo statistics
```

