



Segment Routing uLoop Avoidance

- [Prerequisites for Segment Routing uLoop Avoidance, on page 1](#)
- [Restrictions for Segment Routing uLoop Avoidance, on page 1](#)
- [Information About Segment Routing uLoop Avoidance, on page 1](#)
- [How to Enable Segment Routing uLoop Avoidance, on page 5](#)

Prerequisites for Segment Routing uLoop Avoidance

- The SR: uLoop Avoidance feature will work only if the Topology-Independent Loop-Free Alternate (TI-LFA) feature is configured.

Restrictions for Segment Routing uLoop Avoidance

- Segment routing uLoop avoidance feature does not support double fault at the same time.
- Node up/down events are not supported in OSPF uLoop avoidance.

Information About Segment Routing uLoop Avoidance

Microloops

When changes occur in a network topology because of the failure or restoration of a link or a network device, IP Fast Reroute enables rapid network convergence by moving traffic to precomputed backup paths until regular convergence mechanisms move traffic to a newly computed best path, also known as a post-convergence path. This network convergence may cause short microloops between two directly or indirectly connected devices in the topology. Microloops are caused when different nodes in the network calculate alternate paths at different times and independently of each other. For instance, if a node converges and sends traffic to a neighbor node, which has not converged yet, traffic may loop between the two nodes.


Microloops may or may not result in traffic loss. If the duration of a microloop is short, that is the network converges quickly, packets may loop for a short duration before their time-to-live (TTL) expires. Eventually, the packets get forwarded to the destination. If the duration of the microloop is long, that is one of the routers

in the network is slow to converge, packets may expire their TTL or the packet rate may exceed the bandwidth, or the packets might be out of order, and packets may get dropped.

Microloops that are formed between a failed device and its neighbors are called local uloops, whereas microloops that are formed between devices that are multiple hops away are called remote uloops. Local uloops are usually seen in networks where local loop-free alternate (LFA) path is not available. In such networks, remote LFAs provide backup paths for the network.

The information discussed above can be illustrated with the help of an example topology as shown in the following figure.

Figure 1: Microloop Example Topology




The assumptions in this example are as follows:

- The default metrics is 10 for each link except for the link between Node 3 and Node 6, which has a metric of 50. The order of convergence with SPF backoff delays on each node is as follows:
 - Node 3—50 milliseconds
 - Node 1—500 milliseconds
 - Node 2—1 second
 - Node 2—1.5 seconds

A packet sent from Node 3 to Node 9, the destination, traverses via Node 6.


If a link is established between Node 6 and Node 7, the shortest path for a packet from Node 3 to Node 9 would be Node 1, Node 2, Node 7, and Node 6 before the packet reaches the destination, Node 9.

Figure 2: Microloop Example Topology—Shortest Path




The following figure shows the Forwarding Information Base (FIB) table in each node before the link between Node 6 and Node 7 is established. The FIB entry contains the prefix of the destination node (Node 9) and the next hop.

Figure 3: Microloop Example Topology—FIB Entry



When the link between Node 6 and Node 7 comes up, microloops occur for the links based on the order of convergence of each node. In this example, Node 3 converges first with Node 1 resulting in a microloop between Node 3 and Node 1. Then, Node 1 converges next resulting in a microloop between Node 1 and Node 2. Next, Node 2 converges next resulting in a microloop between Node 2 and Node 7. Finally, Node 7 converges resolving the microloop and the packet reaches the destination Node 9, as shown in the following figure.

Figure 4: Microloop Example Topology—Microloops




Adding the SPF convergence delay, microloop results in a loss of connectivity for 1.5 seconds, which is the convergence duration specified for node 7.


How Segment Routing Prevents Microloops?

Using the example used to explain microloops, this section explains how segment routing prevents microloops. Node 3 in the example is enabled with the **microloop avoidance segment-routing** command. After the link between Node 6 and Node 7 comes up, Node 3 computes a new microloop on the network.

Figure 5: Microloop Example Topology—Segment Routing



Instead of updating the FIB table, Node 3 builds a dynamic loop-free alternate (LFA) SR path for the destination (Node 9) using a list of segments IDs, which include the prefix segment ID (SID) of Node 7, which is 16007, and the adjacency segment ID (SID) of Node 6, which is 24076.



So, the SR path enables a packet from Node 3 reaches its destination Node 9, without the risk of microloop until the network converges. Finally, Node 3 updates the FIB for the new path.

Use the protected keyword with the **microloop avoidance segment-routing** command, to enable microloop avoidance for protected prefixes only. The **microloop avoidance rib-update-delay milliseconds** command can be used to configure the delay in milliseconds for a node to wait before updating the node's forwarding table and stop using the microloop avoidance policy. The default value for the RIB delay is 5000 milliseconds.

How to Enable Segment Routing uLoop Avoidance

Enabling ISIS - Microloop Avoidance

The following is a sample configuration code snippet to enable microloop avoidance in ISIS.

```
router isis 1
 net 49.0001.0000.0000.0209.00
 metric-style wide
 nsf cisco
 distribute link-state
 segment-routing mpls
 segment-routing prefix-sid-map advertise-local
 fast-reroute per-prefix level-1 all
 fast-reroute per-prefix level-2 all
 fast-reroute ti-lfa level-2
 microloop avoidance segment-routing
 bfd all-interfaces
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2
 !
```

Verifying ISIS - Microloop Avoidance

Use the **show isis private** command to check if the repair path exists or not.

```

Router# show isis private
Process: pdb 0x84BD34E0, name isis, tag 1, index -1, mask 0, lvl 3, handle 84BD34E0
Mode: ipv4 active-ip
idbqueue (4): Te0/0/27 Te0/0/26 Lo1 Lo0
passive_idbq (0):
IPv4:
  ip_pdb 0x84BFFE18, index 3, mask 0x8
  ipv6_idbq (0):
TID 0: topoid 0x00000000 (IPv4)
  num_L2 1, redist_L2 0
  ip_ipdb 0x84C00030
  ip_idbq (4): Lo0 Lo1 Te0/0/26 Te0/0/27

```

ISIS Microloop Avoidance Info:

enable state: enabled for segment-routing microloop avoidance

user config: microloop avoidance segment-routing

delay: 5000,

primary_update_complete_pending: 0

L1:

```

  local_down_event[1]:
  uloop_event_if[1]:
  uloop_spf_on[1]:0, uloop_delay_reqed[1]:0,
  uloop_rt_track_requested[1]:0
  uloop_delay_timer[1] is not running
  uloop_lsp_nbr[1]:
  Saved local neighbor list:

```

L2:

```

  local_down_event[2]:
  uloop_event_if[2]:
  uloop_spf_on[2]:0, uloop_delay_reqed[2]:0,
  uloop_rt_track_requested[2]:0
  uloop_delay_timer[2] is not running
  uloop_lsp_nbr[2]:
  Saved local neighbor list:
  ISIS Microloop avoidance L(1) lost adjacency queue is empty
  ISIS Microloop avoidance L(2) lost adjacency queue is empty

```

Enabling OSPF - Microloop Avoidance

The following is a sample configuration code snippet to enable microloop avoidance in ISIS.

```

router ospf 1
nsr
nsf ietf
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
fast-reroute per-prefix ti-lfa area 0
microloop avoidance segment-routing
network 170.50.0.0 0.0.255.255 area 0
network 203.203.203.203 0.0.0.0 area 0
bfd all-interfaces
mpls traffic-eng router-id Loopback1
mpls traffic-eng area 0
!

```

Verifying OSPF - Microloop Avoidance

Use the `show ip ospf segment-routing microloop-avoidance` command to check if the repair path exists or not.

```
Router#show ip ospf segment-routing microloop-avoidance
```

```
OSPF Router with ID (213.213.213.213) (Process ID 1)
```

```
SR Microloop Avoidance is configured, delay 5000 msec
```

```
Area with ID (0)
```

```
Base Topology (MTID 0)
```

```
SR Microloop Avoidance is enabled and not running
```

```
Last topology change details:
```

```
Near end: 213.213.213.213
```

```
Far end: 215.215.215.215
```

```
Event: Link Up
```

