

Troubleshooting Cisco Express Forwarding Routing Loops

Document ID: 26083

- Introduction**
- Prerequisites**
 - Requirements
 - Components Used
 - Conventions
- Network Diagram**
- Problem**
- Troubleshoot**
- Solution**
- Related Information**

Introduction

This document helps troubleshoot Cisco Express Forwarding (CEF) routing loops and sub-optimal routing caused by a valid cached Cisco Express Forwarding adjacency that points out the incorrect interface. An adjacency with an incorrect interface is created because of these reasons:

- A static route points directly to a multi-access interface.
- A valid Cisco Express Forwarding adjacency is built as a result of Proxy Address Resolution Protocol (ARP) replies.

Prerequisites

Requirements

Use these resources in order to better understand some of the concepts this document uses:

- Cisco Express Forwarding Overview
- Route Selection in Cisco Routers

Components Used


This document is not restricted to specific software and hardware versions.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Network Diagram

Router R1 connects to R3 via Serial 8/0, and router R2 connects to R4 via Serial 8/0. R1 and R2 are connected via Ethernet 0/0, as this figure shows.



- R2 receives external Border Gateway Protocol (eBGP) prefix updates for 10.10.34.0/24 from R4. R2 propagates this prefix to R1 via internal BGP (iBGP).
- R2 has a static default route (0.0.0.0/0) that points to R4's Serial 8/0 IP address 10.10.24.4.
- R2 also has a back-up floating default route (IP route 0.0.0.0 0.0.0.0 Ethernet0/0 10) that points to interface Ethernet 0/0 to route packets if the serial connection between R2 and R4 fails.
- R1 has a default route that points to R3's Serial 8/0 with the IP address 10.10.13.3.

Problem

IP traffic destined for 10.10.34.0/24 gets looped between R1 and R2. Observe the **tracert** command output on R1.

```
R1#tracert 10.10.34.4

Type escape sequence to abort.
Tracing the route to 10.10.34.4

 0 192.168.12.2 20 msec 20 msec 20 msec
 1 192.168.12.1 8 msec 12 msec 8 msec
 2 192.168.12.2 8 msec 8 msec 12 msec
 3 192.168.12.1 12 msec ...
```

Note that traffic destined for 10.10.34.4 hops between R1's Ethernet 0/0 (IP address 192.168.12.1) and R2's Ethernet 0/0 (IP address 192.168.12.2). Ideally, traffic from R1 destined for 10.10.34.0/24 needs to go to R2 because of the iBGP learned prefix 10.10.34.0/24. Then, from R2, the traffic should route to R4. However, the **tracert** command output confirms a routing loop between R1 and R2.

R1
<pre>hostname R1 ! ip subnet-zero ! ip cef !</pre>

```

interface Ethernet0/0
 ip address 192.168.12.1 255.255.255.0
!
interface Serial8/0
 ip address 10.10.13.1 255.255.255.0
!
router bgp 11
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.10.13.3 remote-as 12
 neighbor 192.168.12.2 remote-as 11
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.10.13.3

```

R2

```

hostname R2
!
ip cef
!
interface Ethernet0/0
 ip address 192.168.12.2 255.255.255.0
!
interface Serial8/0
 ip address 10.10.24.2 255.255.255.0
!
router bgp 11
 no synchronization
 bgp log-neighbor-changes
 network 192.168.12.0
 neighbor 10.10.24.4 remote-as 10
 neighbor 192.168.12.1 remote-as 11
 neighbor 192.168.12.1 next-hop-self
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.10.24.4
ip route 0.0.0.0 0.0.0.0 Ethernet0/0 10
!

```

Troubleshoot

Since the packets destined for 10.10.34.4 get looped between R1 and R2, start to troubleshoot. First check the IP routing on R1. The **show ip route 10.10.34.0** command output confirms the next hop of 192.168.12.2 for packets destined to 10.10.34.0/24. This matches with the **traceroute** command first hop, where packets are sent to next hop 192.168.12.2, which confirms that packets are switched correctly on R1.

```

R1#show ip route 10.10.34.0
Routing entry for 10.10.34.0/24
  Known via "bgp 11", distance 200, metric 0
  Tag 10, type internal
  Last update from 192.168.12.2 00:22:59 ago
  Routing Descriptor Blocks:
    * 192.168.12.2, from 192.168.12.2, 00:22:59 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1

```

The next step is to check the IP routing table of R2. As this **show ip route 10.10.34.0** command output shows, packets destined to 10.10.34.0 should be routed out to next hop 10.10.24.4 on Serial 8/0. However, the **traceroute** command shows packets switched back to R1 to the IP address 192.168.12.1. Further investigation is needed into why packets destined to 10.10.34.0 are switched on R2 to next hop 192.168.12.1 (as in the

output of the **traceroute** command) instead of to 10.10.24.4.

```
R2#show ip route 10.10.34.0
Routing entry for 10.10.34.0/24
  Known via "bgp 11", distance 20, metric 0
  Tag 10, type external
  Last update from 10.10.24.4 00:42:32 ago
  Routing Descriptor Blocks:
  * 10.10.24.4, from 10.10.24.4, 00:42:32 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

At this point it is important to understand that in a Cisco Express Forwarding–switched network, a packet forwarding decision consists of:

- A routing table lookup for the longest prefix match.
- A forwarding information base (FIB) lookup.

Since the routing table is verified, look at the Cisco Express Forwarding FIB. In the results of the **show ip cef 10.10.34.4 detail** command, note that Cisco Express Forwarding switches 10.10.34.4 out Ethernet 0/0 instead of next hop 10.10.24.4 out Serial 8/0 (as shown in the **show ip route 10.10.34.0** command output). This discrepancy creates loops in the network.

```
R2#show ip cef 10.10.34.4 detail
10.10.34.4/32, version 19, cached adjacency 10.10.34.4
0 packets, 0 bytes
  via 10.10.34.4, Ethernet0/0, 0 dependencies
    next hop 10.10.34.4, Ethernet0/0
    valid cached adjacency
```

The next step is to look at the Cisco Express Forwarding adjacency table and see how Cisco Express Forwarding learns to switch packets out Ethernet 0/0. Notice the adjacency is built because of ARP.

```
R2#show adjacency ethernet 0/0 detail | begin 10.10.34.4
IP           Ethernet0/0           10.10.34.4(5)
                    50 packets, 2100 bytes
                    AABCC006500AABCC0066000800
                    ARP           03:02:00
```

This **show ip arp** command output is confirmation.

```
R2#show ip arp 10.10.34.4
Protocol Address          Age (min)  Hardware Addr  Type   Interface
Internet 10.10.34.4        60         aabb.cc00.6500 ARPA   Ethernet0/0
```

Next, find out why this ARP entry was created when there is an IP route in the routing table. Look at the routing table again.

```
R2#show run | include ip route 0.0.0.0
ip route 0.0.0.0 0.0.0.0 10.10.24.4
ip route 0.0.0.0 0.0.0.0 Ethernet0/0 10
```

If the serial connection fails between R2 and R4, all traffic is routed with the use of a floating static route out Ethernet 0/0 because R2 has a floating static route that points to the multi–access interface Ethernet 0/0, and not to the Ethernet IP address 192.168.12.1 of R1. Therefore, for all unknown destinations, Router R2 sends out an ARP request through the Ethernet0/0 interface. In this case, R2 has lost the more specific route to the 10.10.34.0 network. Therefore, when the data packet arrives for the hosts on this network, it generates an ARP request via the Ethernet interface. Since Proxy ARP is enabled by default on R1's Ethernet interface and it has a default route that points to R3, it responds back with an Proxy ARP reply with its own MAC address.

Hence, R2 sends all traffic to R1, and R1 forwards all traffic with the use of its default route (0.0.0.0/0) out to AS 12, and consequently to 10.10.34.4 via the Internet.

When R2 receives the proxy ARP reply from R1, it creates a /32 valid Cisco Express Forwarding adjacency that points out interface Ethernet 0/0. This Cisco Express Forwarding entry does not age out until the proxy ARP router R1 is present on the Ethernet segment. Thus, the /32 Cisco Express Forwarding entry continues to be used to Cisco Express Forwarding—switch the packets, even after the serial connection between R2 and R4 is back up and the routing table default route points out Serial 8/0 towards AS 10. The result is a routing loop.

Finally, look at the logs and see if the serial link (s8/0) flapped. This causes a floating static route to be installed in the routing table which then leads to proxy ARP and results in the installation of a Cisco Express Forwarding entry of 10.10.34.4/32 in the Cisco Express Forwarding FIB.

```
R2#show log | beg Ethernet0/0
[...]
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial8/0, changed state to down
%BGP-5-ADJCHANGE: neighbor 10.10.24.4 Down Interface flap
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial8/0, changed state to up
%BGP-5-ADJCHANGE: neighbor 10.10.24.4 Up
```

The logs confirm the cause. In summary, these steps show the sequence of events:

1. Serial 8/0 on R2 goes down.
2. R2 has a packet destined to 10.10.34.4.
3. R2 follows the backup default route pointed directly to the Ethernet 0/0.
4. R2 sends an ARP request for 10.10.34.4.
5. R1 (Proxy) replies to the ARP request with its own MAC address to R2.
6. R2 now has an ARP entry for 10.10.34.4 with the MAC address of R1.
7. R2 creates a Cisco Express Forwarding adjacency for 10.10.34.4, and a 10.10.34.4/32 entry is installed in the Cisco Express Forwarding table (FIB) for this destination via Ethernet 0/0. This Cisco Express Forwarding entry is maintained for as long as the ARP entry is valid or until R1 is present on the Ethernet segment.
8. Serial 8/0 on R2 comes up.
9. R2 learns eBGP route 10.10.34.0/24 from R4 with next hop 10.10.24.4 and installs the route in the IP routing table.
10. R1 learns prefix 10.10.34.0/24 via iBGP from R2 and installs it in the IP routing table.
11. R1 has a packet destined for 10.10.34.4.
12. R1 looks into its routing table, matches iBGP prefix routes to R2, and routes to R2.
13. R2 receives a packet destined for 10.10.34.4. Since it already has a Cisco Express Forwarding entry for 10.10.34.4/32 that points to Ethernet 0/0 in its FIB table with the MAC address of R1, it sends the packet back to R1 without looking at the routing table. This creates a loop.

Solution

Replace the floating static route that points directly to the Ethernet 0/0 with one that points to a next hop address.

```
R2(config)#no ip route 0.0.0.0 0.0.0.0 ethernet 0/0 10
R2(config)# ip route 0.0.0.0 0.0.0.0 192.168.12.1 10
```

When you have a static route that points to the next hop IP address instead of a multi-access interface Ethernet 0/0, it stops R2 from sending ARP requests for all destinations. The packets are routed and switched based on the next hop 192.168.12.1. Therefore, any ARP Cisco Express Forwarding entries and loops are avoided.

Observe the Cisco Express Forwarding entry on R2 that points to the correct interface Serial 8/0.

```
R2#show ip cef 10.10.34.4
10.10.34.0/24, version 32, cached adjacency to Serial8/0
0 packets, 0 bytes
  via 10.10.24.4, 0 dependencies, recursive
    next hop 10.10.24.4, Serial18/0 via 10.10.24.0/24
    valid cached adjacency
```

Related Information

- [Troubleshooting Load Balancing Over Parallel Links Using Cisco Express Forwarding](#)
 - [How to Verify Cisco Express Forwarding Switching](#)
 - [Troubleshooting Prefix Inconsistencies with Cisco Express Forwarding](#)
 - [Troubleshooting Incomplete Adjacencies with Cisco Express Forwarding](#)
 - [Cisco Express Forwarding Support Page](#)
 - [IP Routing Protocols Support Pages](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2009 – 2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 26083
