

Internet Routing Table Growth Causes %ROUTING-FIB-4-RSRC_LOW Message on Trident-Based Line Cards



Document ID: 116999

Contributed by Jean-Christophe Rode, Eddie Chami, and David Powers,
Cisco TAC Engineers.
Jan 06, 2014

Contents

Introduction

Problem

Command Output to Capture

Solution

Related Information

Introduction

This document describes how to identify and resolve a common problem caused by growth of the Internet routing table: a Trident-based line card reaches its prefix limit, the message %ROUTING-FIB-4-RSRC_LOW occurs, and there is traffic loss on the line cards.

Problem

As the Internet routing table approaches 500,000 prefixes, problems might occur on a Cisco ASR 9000 Series Aggregation Services Router with Trident-based (Ethernet) line cards that use the default scale profile. An ASR 9000 Trident-based line card can support a maximum of 512,000 Layer 3 (L3) prefixes by default. This limit can easily be reached when the router carries the full Internet table as well as Interior Gateway Protocol (IGP) routes and L3 VPN routes.

Typhoon-based (Enhanced Ethernet) line cards support more prefixes by default, so they have more capacity and generally do not require tuning. A Typhoon-based line card supports four million IPv4 and two million IPv6 prefixes by default.

See ASR 9000 Series Line Card Types for an explanation of the differences between Trident-based and Typhoon-based line cards.

Note: Sudden increases in the number of Internet prefixes can happen occasionally. Even if the current Internet table has room for a few thousand more prefixes before it reaches the default limit of 512,000 prefixes on Trident-based line cards, a sudden burst of prefixes on the Internet can result in an outage on these line cards.

The router logs messages such as these when the limit for a Trident-based line card is reached:

```
LC/0/2/CPU0:Dec 6 01:24:14.110 : fib_mgr[169]: %ROUTING-FIB-4-RSRC_LOW :  
CEF running low on DATA_TYPE_TABLE_SET resource memory. CEF will now begin  
resource constrained forwarding. Only route deletes will be handled in this  
state, which may result in mismatch between RIB/CEF. Traffic loss on certain  
prefixes can be expected. CEF will automatically resume normal operation, once
```

the resource utilization returns to normal level.

Once the Trident-based line cards begin to display the %ROUTING-FIB-4-RSRC_LOW message, an outage for some prefixes occurs. Cisco recommends that you proactively review and plan for this problem because there is not always an easy solution after the problem occurs.

Command Output to Capture

Notes:

Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

The Output Interpreter Tool (registered customers only) supports certain *show* commands. Use the Output Interpreter Tool in order to view an analysis of *show* command output.

Capture output from these commands in order to analyze the problem:

- *term length 0*
- *show install active summary*
- *show platform*
- *show running-config*
- *show route vrf all afi-all safi-all sum*
- *show mpls forwarding summary*
- *show hw-module profile scale*
- *show mpls forwarding summary*
- *show cef vrf all summary*
- *show cef resource location* location of Trident line card that reports the message
- *show cef platform resource location* location of Trident line card that reports the message
- *show cef platform resource summary location* location of Trident line card that reports the message (Cisco IOS® XR Software Release 4.3.2, 5.1.1, and later)

The *show cef platform resource location* command provides the number of entries for each hardware resource and the corresponding maximum number of entries.

```
RP/0/RSP0/CPU0:router#sh cef platform resource location 0/1/CPU0
```

```
Node: 0/1/CPU0
```

```
<snip>
```

```
IPV4_LEAF_P usage is same on all NPs
```

```
NP: 0 struct 23: IPV4_LEAF_P (maps to ucode stru = 54)
```

```
Used Entries: 471589 Max Entries: 524288
```

The line card in this example carries 471,000 prefixes, which is close to the supported default limit of 512,000 prefixes for Trident-based line cards. In the case of instability (such as convergence or a sudden burst of prefixes on the Internet), the threshold might be crossed, and the line card would enter the out-of-resources mode.

In Cisco IOS Software Release 4.3.2 and later, the *show cef platform resource location* command takes a long time (up to 15 minutes) to complete, so you might conclude that the command is not working. In Release

4.3.2, 5.1.1, and later, use the *show cef platform resource summary location* command instead.

```
RP/0/RSP0/CPU0:router2#show cef platform resource summary loc 0/2/cpu0
```

OBJECT	USED	MAX	AVAILABLE
RPF_STRICT	0	262144	262144
IPv4_LEAF_P	114	4194304	4194190
IPv6_LEAF_P	57	2097152	2097095
LEAF	716	4194304	4193588
TX_ADJ	652	524288	523636
NR_LDI	715	2097152	2096437
TE_NH_ADJ	0	65536	65536
RX_ADJ	27	131072	131045
R_LDI	662	131072	130410
L2VPN_LDI	0	32768	32768
EXT_LSPA	630	524288	523658
IPv6_LL_LEAF_P	0	262144	262144

Solution

A scale profile is a user-configurable setting that tunes the router so it performs more efficiently, dependent on how the router is being used.

- Before you deploy the router to production, configure a scale profile that allows for the current size of the Internet routing table as well as growth and unexpected increases in prefixes.
- Configure a non-default scale profile if possible. If the Trident-based line card does not have too many Layer 2 (L2) VPN forwarding entries, you can configure the L3 scale profile or L3 XL scale profile in order to allocate more resources to L3 forwarding entries.
- The L3 scale profile can handle one million L3 prefixes, which should be enough for the Internet routing table. When Virtual Routing and Forwarding (VRF) tables are used, it may be necessary to increase the limit to 1.3 million with the L3 XL scale profile.
- Decrease the number of routes handled by the router through summarization. This may not be practical.

For more information, see *Configuring Profiles on the Cisco ASR 9000 Series Router*.

The number of L2 VPN forwarding entries (mac-address-table, bridge-domains, and so forth) decreases when the scale profile is changed. This solution should be carefully evaluated when the router provides both L3 and L2 services, because the forwarding resources must be shared between these features.

Feature	Profile		
	default	L3	L3XL
ipv4 prefixes 1D	512,000	1,000,000	1,300,000
ipv6 prefixes 1D	128,000	500,000	650,000
per-subtree prefixes (pfx/VRF)	128,000	128,000	256,000
adjacencies (ARP entries)	128,000/LC	128,000/LC	256k+ /LC
IGP routes	20,000+	50,000+	50,000+
IGP paths for ECMP	32	32	32
BGP paths for ECMP	4	4	4
LAG bundle members	64	64	64
MAC addresses	512,000	128,000	32,000
Bridge domains	8,000	8,000	2,000
EFPs (I2transport subinterfaces)	64,000	64,000	64,000

The ASR9000/XR Understanding Route scale document in the Cisco support forum provides additional useful information.

Use the *hw-module profile scale* command in order to configure the scale profile from the administration configuration mode. If a scale profile is also configured in the global configuration, you should duplicate the configuration in the administration configuration and remove the global configuration.

This example changes the scale profile to the L3 scale profile:

```
RP/0/RSP1/CPU0:router#admin
RP/0/RSP1/CPU0:router(admin)#config
RP/0/RSP1/CPU0:router(admin-config)#hw-module profile scale ?
  default  Default scale profile
  l3       L3 scale profile
  l3xl     L3 XL scale profile
RP/0/RSP1/CPU0:router(admin-config)#hw-module profile scale l3
In order to activate this new memory resource profile, you must manually reboot
the line cards.
RP/0/RSP1/CPU0:router(admin-config)#commit
RP/0/RSP1/CPU0:router(admin-config)#end
RP/0/RSP1/CPU0:router(admin)#exit
RP/0/RSP1/CPU0:router#
```

In order to activate the new profile, the line card must be manually reloaded, which will interrupt traffic through the line card for a few minutes:

```
RP/0/RSP1/CPU0:router#hw-module location 0/0/CPU0 reload
WARNING: This will take the requested node out of service.
Do you wish to continue?[confirm(y/n)]y
RP/0/RSP1/CPU0:router#
```

In very rare cases, there may not be a scale profile that provides the needed number of L2 and L3 forwarding entries. In those cases, the only solution is to upgrade from Trident-based line cards to Typhoon-based line cards, which support four million IPv4 forwarding entries by default.

In a future release, the default scale profile will be changed. Cisco Bug ID CSCu197045 , "Make the layer 3 scale profile the default for Trident linecards," is a feature request that will change the default scale profile to match the current L3 profile and will introduce a new L2 scale profile that matches the current default.

Related Information

- *Configuring Profiles on the Cisco ASR 9000 Series Router*
- *ASR 9000 Series Line Card Types*
- *Technical Support & Documentation – Cisco Systems*

Updated: Jan 06, 2014

Document ID: 116999
