Understand IPv4 Hardware Resources on Catalyst 9000 Switches

Contents

Introduction

This document describes how to understand and verify IPv4 Forwarding Information Base (FIB) hardware usage on Catalyst 9000 series switches.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Catalyst 9200, 9300, 9400, 9500 (Non-High Performance) series switches on Cisco IOS® XE 16.x & 17.x software
- Cisco Catalyst 9500 (High Performance), 9600 series switches on Cisco IOS® XE 16.x & 17.x software

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

When you troubleshoot a feature that does not operate as expected a good reference point is to confirm that hardware is not beyond the scale of the switch in question. While switches can vary in the size of these hardware tables, the verification and troubleshoot methodology remains mostly the same.

This page is also a reference page where you can find information on various features and how to check hardware scale.

Examples are provided for these software versions and device types:

- 16.x & 17.x software trains, as the CLI outputs are significantly different
- Information specific to the 9500 (High Performance) & 9600 as these models allocate Hash and Ternary Content Addressable Memory (TCAM) differently than other Catalyst 9000 switches

This document helps with:

- Understand which table (hash/TCAM) is consumed
- Understand what consumes the table in question
- Understand why a certain table was used (hash versus TCAM)
- Understand logs or error messages that indicate a resource issue
- Actions to take to resolve hardware resource allocation issues

---

> **Note**: When switches use Virtual Routing and Forwarding (VRF)s the total usage needs to include consumption in each VRF.

Video

This video covers troubleshooting of FIB hardware resource issues: Troubleshoot Catalyst 9000 Switch FIB Hardware Resource Exhaustion

Terminology

ADJ	Adjacency (table)	Stores next hop information used for packet rewrite
DI	Destination Index	Index that points to the outbound interface
EM	Exact Match	An entry in Hash memory that is a 1:1 match (host route, Directly Connected host)
FIB	Forwarding Information Base	Simplified table with prefixes added by the Routing Information Base (RIB) and Address Resolution Protocol (ARP) tables with a pointer to the ADJ table
FED	Forward Engine	The Application Specific Integrated Circuit (ASIC) (hardware) layer

	Driver	
FMAN-FP	Forward Manager-Forwarding Plane	FMAN-FP manages software objects that add, delete, or modify FED information
LPM	Longest Prefix Match	Any route that is /31 or shorter (/32 routes are EM type)
RI	Rewrite Index	MAC address rewrite information for layer 3 forwarding to the next hop adjacency
RIB	Routing Information Base	The routing table seen in "**show ip route**"
SDM	Switch Database Manager	Software process which allocates switch hardware resources to various features that require them (MAC Addresses, Routes, Access-List Entries)
SI	Station Index	Station Index = packet rewrite information (RI = Rewrite Index) & outbound interface information (DI = Destination Index)
TCAM	Ternary Content-Addressable Memory	A type of memory that stores and queries entries with three different inputs: 0, 1 and X. This type of memory must be used in cases where there can be multiple matches to the same entry, and the resulting Hash for each would not be unique. **This table includes a mask or "X" value that allows it to know if it matches or does not match this entry.**
UADP	Cisco Unified Accessâ„¢ Data Plane	The ASIC architecture used in the switch
Directly Connected	Directly Connected Route	A locally connected host prefix (ARP adjacent)
Indirectly Connected	Indirectly Connected Route	A route that is via a remote next hop to reach
SGT		
SXP		

CTS (Trustsec)		

General Hardware Validation Commands

These commands show high level usage statistics for Hash, TCAM, Interface, and Rewrite resources used. These resources are related, and exhaustion of one of the resources mentioned can affect the ability to fully use other available resources.

Example: A switch can have available Hash / TCAM, but ran out of Adjacencies. The ability to forward packets can be impacted to a destination prefix because the switch cannot program a new rewrite entry.

<#root>

```
show platform hardware fed switch active fwd-asic resource tcam utilization


<-- Hash & TCAM
show platform hardware fed switch active fwd-asic resource utilization


<-- SI/RI/DI/etc (other related resources)


show platform hardware fed switch active fwd-asic resource rewrite utilization


<-- IP Adjacency. LISP adjacency, Tunnel Adjacency, etc
```

> **Note**: See "Chapter: Configuring SDM Templates" in the "System Management Configuration Guide" for more details about these commands

---

IPv4 Scale Syslogs

This scenario demonstrates how each table is used, and what to do if one or another table is at or beyond scale. It also covers dependent resources required to forward to an IP destination.

Symptom: The resource is beyond scale

1. Device or prefix reachability issues. While routes that exist or devices can remain reachable, any new or updated prefixes are not reachable.
2. Log messages indicate the hardware is not able to take new object updates
3. The object layer, which programs software into hardware become congested
4. Absent entries at the impacted hardware layer (in this case the FIB is the impacted layer).

If you run out of a particular IPv4 FIB or Adjacency resource SYSLOG message are generated by the system

IPv4 FIB Log Message	Definition	Recovery Action
%FED_L3_ERRMSG-3-RSRC_ERR: Switch 1 R0/0: fed:**Failed to allocate hardware resource for fib entry**due to hardware resource exhaustion	Hardware reserved for IPv4 FIB entries has run out of space (EM or TCAM)	Summarize routes or take some other action to reduce the scale of FIB entries (this can be EM or TCAM, whichever one is exhausted).
%FED_L3_ERRMSG-3-RSRC_ERR: R0/0: fed:**Failed to allocate hardware resource for adj entry - rc:1**	The Adjacency table is exhausted. This is the table in hardware where next- hop destination MAC addresses are stored.	Reduce the scale number of directly connected (ARP adjacent) hosts

Usage Examples

Usage Examples (non-HP 16.12.x)

Software	Hardware
16.12.5	Catalyst 9200 9300 9400 9500 (Non-High Performance) switches

Baseline Resource Usage

<#root>

Baseline Setup & Usage

C9300#

show version | include IOS

Cisco IOS XE Software,

Version 16.12.05

Cisco IOS

Software [Gibraltar],

Catalyst L3 Switch Software (CAT9K_IOSXE)

, Version 16.12.5

, RELEASE SOFTWARE (fc3)

C9300-48U

C9300##

show ip interface brief | exclude unassigned

```
Interface              IP-Address      OK? Method Status                Protocol
```

<...empty...> <-- no Switch Virtual Interface (SVI) or any IP configured

/// TCAM and Hash ///

C9300#

show platform hardware fed switch actitve fwd-asic resource tcam utilization

```
CAM Utilization for ASIC [0]
Table                                              Max Values    Used Values
------------------------------------------------------------------------------
```

Directly or indirectly connected routes

24576/8192

3/19 <-- 3 hash / 19 TCAM is base usage

/// Adjacencies ///

C9300#

show platform software fed switch active ip adj

```
IPV4 Adj entries
dest     if_name         dst_mac       si_hdl      ri_hdl      pd_flags adj_id Last-modified
----     -------         -------       ------      ------      -------- ----- ------------------------
```

<empty>

C9300#

show platform hardware fed switch active fwd-asic resource rewrite utilization

```
Resource Info for ASIC Instance: 0
Rewrite Data                       Allocated      Free
-----------------------------------------------------
```

PHF_EGRESS_destMacAddress

0

32000 <-- Next hop Dest MAC for packet rewrite

/// SI DI RI resources ///

C9300#

show platform hardware fed switch active fwd-asic resource utilization | include RSC_SI_|RSC_RI__

```
Resource Info for ASIC Instance: 0
Resource Name           Allocated     Free
-----------------------------------------
RSC_RI
```

3

```
       57317
```

<-- Rewrite Index

```
RSC_SI
```

521

```
       64847
```

<-- Station Index

Add SVI Vlan 1 IP address with /24 mask

<#root>

ADD SVI IP with /24 mask length

```
C9300(config)#
```

interface vlan 1

```
C9300(config-if)#
```

ip address 10.10.10.1 255.255.255.0

```
C9300#
```

show ip interface brief | inc up
Vlan1   10.10.10.1

```
        YES     manual    up                    up
```

```
C9300#
```

show platform hardware fed switch active fwd-asic resource tcam utilization

```
CAM Utilization for ASIC [0]
Table                                              Max Values      Used Values
------------------------------------------------------------------------------
```

Directly or indirectly connected routes   24576/8192   6/20 <-- usage = +3 hash, +1 TC

```
C9300#
```

show platform software fed switch active ip adj

```
IPV4 Adj entries
```

```
dest         if_name         dst_mac          si_hdl         ri_hdl         pd_flags adj_id  Last-modifi
----         -------         -------          ------         ------         -------- ----- -------------
```

227.0.0.0

Vlan1

0100.5e00.0000

0x7f4880ce37e8 0x7f4880cf3648 0x0 0xf80004b4 2021/02/26 17:48:47.992

<-- 1 Adj created for mcast

C9300#

show platform hardware fed switch active fwd-asic resource rewrite utilization

```
Resource Info for ASIC Instance: 0
Rewrite Data                       Allocated       Free
-----------------------------------------------------
```

PHF_EGRESS_destMacAddress 1 31999 <-- 1 Adj used for mcast

Add 3 EM prefixes (/32 mask)

<#root>

Configuration adds 3 /32 prefixes and uses 3 Hash Entries

```
interface loopback 1
ip address 10.111.111.1 255.255.255.255
```

<-- Local /32 prefix

!

```
ip route 10.111.111.2 255.255.255.255 vlan 1
```

<-- An Indirect EM prefix (same consumption occurs when learnt via routing protocol)

!

```
arp 10.10.10.100 aaaa.bbbb.cccc arpa
```

<-- Static ARP entry in Vlan 1 (same consumption occurs when learnt dynamically)

Usage Result

/32 Loopback creation, /32 Indirect route, ARP entry only use Hash table. No TCAM used.

C9300#

show platform hardware fed switch active fwd-asic resource tcam utilization

```
CAM Utilization for ASIC [0]
Table                                              Max Values   Used Values
--------------------------------------------------------------------------------
 Directly or indirectly connected routes
```

24576/8192 9/20 <-- usage = 3 Hash, 0 TCAM

Usage Examples (non-HP 17.x)

Software	Hardware
17.x	Catalyst 9200 9300 9400 9500 (Non-High Performance) switches

Baseline Resource Usage

<#root>

Baseline Setup & Usage

C9400#

show version | include IOS

Cisco IOS XE Software,

Version 17.03.02a

Cisco IOS Software

[Amsterdam]

, Catalyst L3 Switch Software (CAT9K_IOSXE),

Version 17.3.2a

, RELEASE SOFTWARE (fc5)

C9400-SUP-1

C9400#

show ip interface brief | exclude unassigned

Interface IP-Address OK? Method Status Protocol

<...empty...> <-- no SVI or any IP configured

C9400#

show platform hardware fed active fwd-asic resource tcam utilization

Codes:

EM - Exact_Match

, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC [0]
Table Subtype Dir

Max

Used

%Used

V4

V6 MPLS Other

IP Route Table

EM

I 49152

3

0.01%

2 0 1 0 <-- 3 hash

IP Route Table

TCAM

I 65536

15

0.02%

6 6 2 1 <-- 15 TCAM

C9400#

show platform software fed active ip adj

IPV4 Adj entries
dest if_name dst_mac si_hdl ri_hdl pd_flags adj_id La

```
----            -------                 -------          ------           ------     -------- -----   --
```

<empty>

!!! New CLI displays multiple resources in one place !!!

New CLI collects usage information for all resources related to all IP Routing in one output

C9400#

show platform hardware fed active fwd-asic resource features ip-adjacency utilization

```
IPv4 unicast adjacency resource info
Resource Info for ASIC Instance: 0 [A:0, C:0]
Shared Resource Name
```

Allocated

```
        Free              Usage%
-------------------------------------------------------------------------------
```

RSC_RI

```
                                    3               57317               0.01
```

<-- Rewrite Index

RSC_SI

```
                                520             64848               0.80
```

<-- Station Index

```
Rewrite Data                              Allocated        Free               Usage%
-------------------------------------------------------------------------------
```

PHF_EGRESS_destMacAddress

```
                0               32000               0.00
```

<-- Next hop Dest MAC for packet rewrite

```
CAM Table Utilization Info               Allocated         Free               Usage%
-------------------------------------------------------------------------------
```

IP Route table Host/Network

3/ 15

49149/327

```
      0.01/ 0.05
```

<-- Hash / TCAM Table usage

Add SVI Vlan 1 IP address with /24 mask

<#root>

C9400(config)#

interface vlan 1

C9400(config-if)#

ip address 10.10.10.1 255.255.255.0

C9400#

show ip interface brief | exclude unassigned

```
Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  10.10.10.1      YES manual up                    up
```

C9400#

show platform hardware fed active fwd-asic resource features ip-adjacency utilization

Resource Info for

ASIC Instance: 1

```
 [A:0, C:1]
Shared Resource Name
```

Allocated

```
         Free             Usage%
-------------------------------------------------------------------------------
RSC_RI
```

4

```
               57316              0.01
```

<-- 1 Rewrite Index

```
RSC_SI                                  520               64848              0.80
```

Rewrite Data

Allocated

```
         Free              Usage%
-------------------------------------------------------------------------------
PHF_EGRESS_destMacAddress
```

1

31999

```
              0.00
```

<-- 1 Adj

used for mcast

```
CAM Table Utilization Info
```

Allocated

```
         Free               Usage%
--------------------------------------------------------------------------------
IP Route table Host/Network
```

6/ 16

```
         49146/32752    0.01/ 0.05
```

<-- 3 Hash + 1 TCAM

```
C9400#
```

show platform software fed active ip adj

```
IPV4 Adj entries
dest            if_name                 dst_mac          si_hdl         ri_hdl         pd_flags adj_id
----            -------                 -------          ------         ------         -------- -----
```

227.0.0.0

```
       Vlan1
```

0100.5e00.0000

```
   0x7fd8fd1654c8 0x7fd8fc8e6098 0x0      0xf8000444  2015/01/01 00:49:54.758
```

<-- 1 Adj created for mcast

Add 3 EM prefixes (/32 mask)

<#root>

Configuration adds 3 /32 prefixes and uses 3 Hash Entries

```
interface loopback 1
ip address 10.111.111.1 255.255.255.255
```

<-- Local /32 prefix

```
!
ip route 10.111.111.2 255.255.255.255 vlan 1
```

<-- An Indirect EM prefix (same consumption occurs when learnt via routing protocol)

```
!
arp 10.10.10.100 aaaa.bbbb.cccc arpa
```

<-- Static ARP entry in Vlan 1 (same consumption occurs when learnt dynamically)

Usage Result

```
/32 Loopback creation, /32 Indirect route, ARP entry only use Hash table. No TCAM used.
```

```
C9400#
show platform hardware fed active fwd-asic resource tcam utilization

Codes: EM - Exact_Match, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC [0]
Table                   Subtype       Dir      Max
Used
    %Used
V4
       V6    MPLS    Other
--------------------------------------------------------------------------------------------------------
IP Route Table          EM
          I       49152
   9
    0.02%
8
        0       1        0
<-- Previously was 6, + 3 for /32 EM


IP Route Table
TCAM
          I       65536
16
    0.02%       8        6       2        1
<-- Previously was 16, no change


C9400#show platform hardware fed active fwd-asic resource features ip-adjacency utilization

IPv4 unicast adjacency resource info
Resource Info for ASIC Instance: 1
[A:0, C:1] <-- ASIC 0, Core 1


Shared Resource Name
Allocated
        Free              Usage%
--------------------------------------------------------------------------------
RSC_RI                                  5
             57315              0.01
<-- One Rewrite index


RSC_SI                                  522
```

```
           64846               0.80
```

<-- Two Station Index

```
Rewrite Data
```

Allocated

```
        Free               Usage%
--------------------------------------------------------------------------------
```

PHF_EGRESS_destMacAddress   2

```
                31998               0.01
```

<-- One Dest MAC used for ARP entry

```
CAM Table Utilization Info              Allocated        Free                Usage%
--------------------------------------------------------------------------------
IP Route table Host/Network
```

9/ 16

```
        49143/32752     0.02/ 0.05
```

<-- 9 EM, 16 TCAM

Usage Examples (HP & 9600 17.x)

Software	Hardware
17.x	Catalyst 9500(High Performance), 9600 switches

> **Note**: For 9500(High Performance) and 9600 only 17.x CLIs are given. Consult the previous section for 16.x examples.

Baseline Resource Usage

<#root>

Baseline Setup & Usage

9500H#

show version | include IOS

```
Cisco IOS XE Software, Version 17.04.01
Cisco IOS Software [Bengaluru], Catalyst L3 Switch Software (CAT9K_IOSXE), Version 17.4.1, RELEASE SOFTW
```

C9500-24Y4C

C9500H#

show ip interface brief | exclude unassigned

Interface   IP-Address   OK? Method Status   Protocol

<...empty...> <-- no SVI or any IP configured

C9500H#

show platform hardware fed active fwd-asic resource tcam utilization

Codes:

EM - Exact_Match

, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC [0]
Table   Subtype   Dir

Max

Used

%Used

V4

     V6   MPLS   Other

IP Route Table

EM/LPM

       I   65536

3

     0.00%

2   0   1   0 <-- 3 hash (EM/LPM)

IP Route Table

TCAM

         I   1536

11

0.72%

6 3 2 0 <-- 11 TCAM

C9500H#

show platform software fed active ip adj

IPV4 Adj entries

dest	if_name	dst_mac	si_hdl	ri_hdl	pd_flags	adj_id	La

<empty>

!!! New CLI displays multiple resources in one place !!!

New CLI collects usage information for all resources related to all IP Routing in one output

C9500#

show platform hardware fed active fwd-asic resource features ip-adjacency utilization

IPv4 unicast adjacency resource info
Resource Info for

ASIC Instance: 0

[A:0, C:0]

<-- ASIC 0 Allocation

Shared Resource Name

Allocated

Free Usage%

RSC_RI

3 90085 0.00

<-- Rewrite Index

RSC_SI

517 130397 0.39

<-- Station Index

Rewrite Data Allocated Free Usage%

PHF_EGRESS_destMacAddress

0 98304 0.00

```
<-- Next hop Dest MAC for packet rewrite


CAM Table Utilization Info               Allocated         Free               Usage%
-------------------------------------------------------------------------------------

IP Route table Host/Network

                  4

/  12


98300/1524

     0.00/ 0.78

<-- Hash / TCAM Table usage
```

Troubleshoot

Scale Limit and Remediation (UADP 2.0 switches)

This section covers one use case where TCAM is exhausted for IPv4:

- Prefixes simulated in this example are /24
- Protocol BGP
- Platform 9400
- Cisco IOS XE 17.3.2

Baseline Usage

With the 9400 Sup-1 any prefix that is /31 or shorter consumes TCAM, not EM

- The maximum number of IPv4 prefixes that can be added to TCAM is 65535.

> **Note**: The IP Route EM and TCAM is also used by Multiprotocol Label Switching (MPLS) with MPLS labels added to EM first, then overflow to TCAM if the limit is hit. If your device is a MPLS Provider Edge (PE) and allocates VPNv4 labels, this is subtracted from the total number.

> **Note:** It is also important to know that if EM memory fills up first, it is allowed to overflow into TCAM memory. (If the switch scales beyond 49152 EM entries, TCAM can fill with Connected and /32 routes). **However, the reverse does not happen (if TCAM fills up, it does not overflow to EM)**

```
<#root>

C9407R#

show platform hardware fed active fwd-asic resource tcam utilization

Codes: EM - Exact_Match, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC  [0]
 Table                  Subtype      Dir      Max      Used    %Used       V4       V6     MPLS    Other
```

```
 ---------------------------------------------------------------------------------------------------
 Mac Address Table      EM          I       65536       20    0.03%        0        0        0       20
 Mac Address Table      TCAM        I        1024       21    2.05%        0        0        0       21
 L3 Multicast           EM          I       16384        0    0.00%        0        0        0        0
 L3 Multicast           TCAM        I        1024        9    0.88%        3        6        0        0
 L2 Multicast           EM          I       16384        0    0.00%        0        0        0        0
 L2 Multicast           TCAM        I        1024       11    1.07%        3        8        0        0
 IP Route Table         EM          I       49152       31    0.06%       18        0       13        0

 IP Route Table         TCAM        I       65536       24    0.04%       15        6        2        1
<...snip...>
```

Inject 80000 Prefixes Via BGP

<#root>

C9407R#

show bgp vpnv4 unicast all summary

```
BGP router identifier 10.255.255.255, local AS number 65000
BGP table version is 580445, main routing table version 580445
80003 network entries using 20480768 bytes of memory
80003 path entries using 10880408 bytes of memory
16002/8001 BGP path/bestpath attribute entries using 4864608 bytes of memory
8002 BGP AS-PATH entries using 533708 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 36759516 total bytes of memory
BGP activity 420126/340116 prefixes, 475340/395329 paths, scan interval 60 secs
80009 networks peaked at 04:52:57 Jan 1 2015 UTC (01:02:51.236 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.1.2     4        65001     101   40485   501775    0    0 01:25:44        0

192.168.1.6     4        65002   31330      96   580445    0    0 01:23:30    80003   <-- 80K prefixes i
```

<#root>

C9407R#

show platform hardware fed active fwd-asic resource tcam utilization

```
Codes: EM - Exact_Match, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC  [0]
 Table                  Subtype      Dir      Max      Used    %Used       V4       V6     MPLS    Other
 ---------------------------------------------------------------------------------------------------
 Mac Address Table      EM           I       65536       20    0.03%        0        0        0       20
 Mac Address Table      TCAM         I        1024       21    2.05%        0        0        0       21
 L3 Multicast           EM           I       16384        0    0.00%        0        0        0        0
 L3 Multicast           TCAM         I        1024        9    0.88%        3        6        0        0
 L2 Multicast           EM           I       16384        0    0.00%        0        0        0        0
 L2 Multicast           TCAM         I        1024       11    1.07%        3        8        0        0
 IP Route Table         EM           I       49152       31    0.06%       18        0       13        0

 IP Route Table         TCAM         I       65536    65536  100.00%    65527        6        2        1
```

Log message indicates that FED cannot program a FIB type entry

<#root>

```
%FED_L3_ERRMSG-3-RSRC_ERR: R0/0: fed: Failed to allocate hardware resource for fib entry due to hardware
```

Solution

Prefix summarization is required in order to fix this type of scale problem. The method to summarize is dependent on the protoctol used, how contiguous your subnets are, and your specific environment.

- In this example, BGP is the protocol and summarization techniques can be found on this page: IP Routing Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9500 Switches) - Chapter: Configuring BGP
- Other summarization techniques can be found in the Related Information section of this document. (All of the links provided are from the 9500 configuration pages. Consult other platform guides as necessary).

You can also check prefixes installed in hardware by prefix length to help determine how many of each length, and summarize where necessary.

<#root>

```
C9300-48U#

show platform software fed switch active ip route summary

Total number of v4 fib entries = 1024                        <-- total prefix count
Total number succeeded in hardware = 1024                    <-- total successfully installed in hardware


 Mask-Len 0 :-  Total-count 2  hw-installed count 2
 Mask-Len 4 :-  Total-count 2  hw-installed count 2
 Mask-Len 8 :-  Total-count 4  hw-installed count 4

 Mask-Len 24 :-  Total-count 1000  hw-installed count 1000 <-- breakdown by mask length

 Mask-Len 30 :-  Total-count 2  hw-installed count 2
 Mask-Len 32 :-  Total-count 14  hw-installed count 14
<...snip...>
```

Scale Limit and Remediation (UADP 3.0 switches)

As noted in the document "Understand Hardware Resources on Catalyst 9000 Switches" UADP 3.0 based switches use an optimized FIB memory that combines both Exact Match (EM) and Longest Prefix Match (LPM) entries. What is listed as TCAM memory is only used for overflow, collisions, and other exceptions.

If you observe that TCAM is highly utilized and EM/LPM is not as utilized (not near or at maximum values), this can be due to a network design where there are too many prefix lengths used. EM/LPM can only program a specific number of mask lengths, but the number is not static. The amout of masks that EM/LPM can program varies with SDM template, and varies further with the route scale of your network.

In summary, for UADP 3.0 there is a dynamic balance between SDM configuration, route scale (number of

routes), and prefix length variance (number of unique subnet masks). If IP Route Table TCAM is exhausted before IP Route Table EM/LPM is near maximum, it is possible that there is an issue as a result of SDM configuration, route scale (number of routes), and prefix length variance (number of unique subnet masks).

When a FIB scale limit occurs, a syslog is generated which indicates hardware has failed to program:

<#root>

%FED_L3_ERRMSG-3-RSRC_ERR: R0/0: fed: Failed to allocate hardware resource for fib entry due to hardware

---

> **Note**: The syslog only indicates a generic failure, not specifically EM/LPM, TCAM, or otherwise. You need to run additional commands to see which has an issue.

9500-48Y4C (9500H / High Performance - UADP 3.0 based switch)

Example of TCAM used up before EM/LPM:

- The output shown indicates that IP Route Table (FIB) TCAM is full, while EM/LPM is not.
- Understand that the switch allocates prefixes and correspondent masks to multiple instances of fixed-size tables in EM/LPM memory
- A given table instance represents only one IPv4 or IPv6 prefix length. The mask is not saved for every prefix entry in hardware, only for the table instance as a whole, which saves valuable hardware space.
- The allocation of prefix masks to hardware tables is dynamic and no input is required for it to operate.

<#root>

Switch#

show platform hardware fed active fwd-asic resource tcam utilization

Codes: EM - Exact_Match, I - Input, O - Output, IO - Input & Output, NA - Not Applicable

CAM Utilization for ASIC [0]

Table	Subtype	Dir	Max	Used	%Used	V4	V6	MPLS	Other
IP Route Table									
EM/LPM									
		I	212992	134345	**63.08%**	2	0	1	0
<-- 63% (EM is not near or at its limit)									
IP Route Table									
TCAM									
		I	1536	1516	**98.70%**	6	6	2	1

<-- 98% (TCAM is used before EM has reached scale limit)

In this scenario, EM/LPM tables can be underutilized:

- For each new table allocated, a specific prefix mask length is dedicated to that table. All entries for that table are now restricted to that mask length. If you have less than the maximum amount of entries for that table instance, effective utilization is reduced by the amount of entries that remain in that table instance. They cannot be reclaimed for some other mask length.
- If you repeat this suboptimal utilization across multiple mask lengths, effective utilization of EM/LPM memory is reduced significantly.
- When EM/LPM tables are no longer available for new mask lengths, those prefixes with that mask are instead installed and overflow into TCAM. Limited TCAM space fills quickly.
- The amount of tables available to EM/LPM varies based on your SDM configuration.

A wide distribution of mask lengths, combined with masks that have a low amount of prefixes, leads to scenarios where many tables are allocated, and many are underutilized. This causes "**Failed to allocate hardware resource for fib**..." errors to appears before ideal maximum hardware utilization is reached on UADP 3.0 based switches.

As the number of table instances allocated to a single prefix length increases, there is the potential to not have enough tables available for prefixes of other lengths.

Sub-optimal usage can occur when the amount of entries of a prefix length does not fully utilize the last table it is part of.

This example illustrates 4500 eBGP prefixes learnt on a 9500-48Y4C switch.

- Each table can have about 2048 routes / entries.
- The hardware has tables of two fixed sizes, 2048 and 8192. They are dynamically allocated based on route scale, prefix usage, and SDM template used
- For this example, three 2048 sized tables are marked as /32 to hold 4500 routes of /32 mask length. Some of the 2048 sized tables have space for /32 IPv4 routes allocated only.

This can be observed through this CLI:

<#root>

```
F241.03.23-C9500-2#
```

show platform software fed active hash l3unicast

```
******************** TABLE INFO ASIC 0 *********************************
Index  Hash Id   Table Id  Size   Used   Mode      Asic   Core   BaseIdx
1      8         0         2048   11     Normal    0      0      0
2      8         1         2048   0      Normal    0      0      2048
...
```

Total Entries: 11 <<<<

minimal entries, no routes yet injected from eBGP

```
************************** MASK INFO *******************************
Mask Id   Table Count  Size   Available Hash Entries  Overflow Entries
```

```
1         0            0      0         0             1
...
33        1            2048   2037      7             0
```

<<<< Mask ID 33 = IPv4 Mask 32. Hardware Masks beyond 33 are for IPv6, MPLS, or other features.

```
...
```

NOTE: 7 hash (EM/LPM) entries are used already, 0 overflow (TCAM)

```
******* MASK ALLOCATION INFO PER TABLE FOR ASIC 0 *******
Asic: 0 Core: 0 Hash Id: 8 Table Id: 0
Index   MASK    GMR/LPM Idx     Used
0       33      2               7
1       162     2               3
2       164     0               1

Asic: 0 Core: 1 Hash Id: 11 Table Id: 7
Index   MASK    GMR/LPM Idx     Used
0       0       0               0
1       139     2               0
```

Inject 4500 eBGP IPv4 routes;

```
F241.03.23-C9500-2#
```

show platform software fed active hash l3unicast

```
******************** TABLE INFO ASIC 0 *********************************
Index  Hash Id   Table Id
```

Size

Used

```
   Mode      Asic   Core   BaseIdx
1      8         0         2048   2048   Normal    0      0      0
```

<<< Table index 1 fully utilized

```
2      8         1         2048   423    Normal    0      0      2048
```

<<< Table index 2 423 / 2048 utilized, must be used for /32 IPv4 pre

```
...
56     11        7         2048   2029   Normal    0      1      112640

Total Entries: 4500

************************** MASK INFO *********************************
Mask Id   Table Count  Size   Available Hash Entries  Overflow Entries
...
33        3            6144   1644      4496          1
```

<<< Hardware Mask 33 increased table count to 3, 1644 "available" spots for /33 (for example IPv4 /32)

```
******* MASK ALLOCATION INFO PER TABLE FOR ASIC 0 *******
```

```
Asic: 0 Core: 0 Hash Id: 8 Table Id: 0
Index   MASK    GMR/LPM Idx     Used
0       33      2               2044
1       162     2               3
```

<<< Another interface is configured with IPv6, hence the/162 mask. /162 mask in hardware = /128 in IPv6.

```
Asic: 0 Core: 0 Hash Id: 8 Table Id: 1
Index   MASK    GMR/LPM Idx     Used
0       33      2               423

Asic: 0 Core: 1 Hash Id: 11 Table Id: 7
Index   MASK    GMR/LPM Idx     Used
0       33      2               2029
/>
```

In the next example, the minimum amount of prefixes are injected to force a new table to be used for every mask-length from /11 to /32.

This is achieved through usage of the "Distribution" SDM Template on Cisco IOS XE 17.3 - which has 56 tables that can be dynamically allocated for FIB entries.

- /32 length is scaled up to 74000 prefixes
- /12 - 300 prefixes
- /11 - 1250 prefixes
- The mask lengths from /13 to /32 contain about 12 prefixes. Any less, and the hardware can instead choose to use TCAM to store those prefixes of a given mask length.

Note: The hardware algorithm which governs EM/LPM and TCAM allocation is complex, which balances both the user configuration and the limitations of the hardware. Results presented in this document are not fully representative of how a system can behave in a production network.

<#root>

```
F241.03.23-C9500-2#
```

show platform hardware fed active fwd-asic resource tcam utilization | include Subtype|IP Route

```
 Table                  Subtype      Dir      Max     Used    %Used       V4       V6     MPLS    Other
 IP Route Table         EM/LPM       I      114688       11    0.01%        7        3        1        0
 IP Route Table         TCAM         I        1536       15    0.98%        7        6        2        0
```

Inject the routes...

***Jan 8 16:17:47.762: %FED_L3_ERRMSG-3-RSRC_ERR: R0/0: fed: Failed to allocate hardware resource for fib**

```
F241.03.23-C9500-2#
```

show platform hardware fed active fwd-asic resource tcam utilization | include Subtype|IP Route

```
Table                  Subtype      Dir      Max     Used    %Used       V4       V6     MPLS    Other
 IP Route Table         EM/LPM       I      114688    73326
```

63.94%

    73322   3   1   0

<<< EM/LPM at 63.94%

 IP Route Table   TCAM   I   1536   1535

99.93%

     1527   6   2   0

<<< TCAM nearly full

F241.03.23-C9500-2#

show platform software fed active ip route summary

Total number of v4 fib entries = 75789
Total number succeeded in hardware = 74847

 Mask-Len 0 :- Total-count 1 hw-installed count 1
 Mask-Len 4 :- Total-count 1 hw-installed count 1
 Mask-Len 8 :- Total-count 2 hw-installed count 2
 Mask-Len 11 :- Total-count 1250 hw-installed count 1250
 Mask-Len 12 :- Total-count 300 hw-installed count 300
 Mask-Len 13 :- Total-count 12 hw-installed count 12
 Mask-Len 14 :- Total-count 12 hw-installed count 12
 Mask-Len 15 :- Total-count 12 hw-installed count 12
 Mask-Len 16 :- Total-count 12 hw-installed count 12
 Mask-Len 17 :- Total-count 12 hw-installed count 12
 Mask-Len 18 :- Total-count 12 hw-installed count 12
 Mask-Len 19 :- Total-count 12 hw-installed count 12
 Mask-Len 20 :- Total-count 12 hw-installed count 12
 Mask-Len 21 :- Total-count 12 hw-installed count 12
 Mask-Len 22 :- Total-count 12 hw-installed count 12
 Mask-Len 23 :- Total-count 12 hw-installed count 12
 Mask-Len 24 :- Total-count 12 hw-installed count 12
 Mask-Len 25 :- Total-count 12 hw-installed count 12
 Mask-Len 26 :- Total-count 12 hw-installed count 12
 Mask-Len 27 :- Total-count 12 hw-installed count 12
 Mask-Len 28 :- Total-count 12 hw-installed count 12
 Mask-Len 29 :- Total-count 12 hw-installed count 12
 Mask-Len 30 :- Total-count 12 hw-installed count 12
 Mask-Len 31 :- Total-count 12 hw-installed count 12
 Mask-Len 32 :- Total-count

74007

  hw-installed count

73065

 <<<

 74007 total /32 known by software, 73065 successfully installed in hardware

F241.03.23-C9500-2#

show platform software fed active hash l3unicast

```
******************** TABLE INFO ASIC 0 *********************************
Index  Hash Id   Table Id  Size
```

Used

```
   Mode      Asic   Core   BaseIdx
1      8         0         2048
```

2048

```
   Normal    0      0      0
...
56     11        7         2048
```

12

```
     Normal    0      1      112640
```

<<< Table indexes 1-56 (varies with SDM) will exist and all show some amount in "used" column

```
************************** MASK INFO *********************************
Mask Id   Table Count  Size   Available Hash Entries  Overflow Entries
1         0            0      0         0             1
5         0            0      0         0             1
9         0            0      0         0             2
10        0            0      0         0             0
11        0            0      0         0             0
12        1            2048   809       1239          11
13        1            2048   1759      289           11
14        1            2048   2047      1             11
15        1            2048   2047      1             11
16        1            2048   2047      1             11
17        1            2048   2036      12            0
18        1            2048   2036      12            0
19        1            2048   2036      12            0
20        1            2048   2036      12            0
21        1            2048   2036      12            0
22        1            2048   2036      12            0
23        1            2048   2036      12            0
24        1            2048   2036      12            0
25        1            2048   2036      12            0
26        1            2048   2047      1             11
27        1            2048   2047      1             11
28        1            2048   2047      1             11
29        1            2048   2047      1             11
30        1            2048   2047      1             11
31        1            2048   2047      1             11
32        1            2048   2047      1             11
33        35
```

71680

```
  0
```

71676

1389

<<< Mask ID /33 allocated 35 tables, 71680 entries, 716676 in EM/LPM (Hash). There are 1389 hw-installed

To resolve this scenario, use one or more of these options in order of priority

- Investigate alternative SDM templates:
  - Increased FIB / IP Route scale templates increase total hardware tables available to this FIB. This can improve effective utilization (used entries / available entries) as well as total scale.
  - Conversely - lower scale SDM templates for FIB / IP Route reduce the total amount of hardware tables available to the FIB.
- Eliminate one or more prefix (mask) lengths entirely from your route table
- Summarize routes (reduce total volume of prefixes)

Note: See "Chapter: Configuring SDM Templates" in the "System Management Configuration Guide" for Catalyst 9500 and Catalyst 9600 Series Switches to learn more about SDM templates.

Scenario: SGT / SXP Mappings | Trustsec Scale

Commands to Collect for TAC

The most common hardware resource problems related to IPv4 utilization are covered in this guide, with appropriate remediation steps. However, in the event that this guide did not resolve your issue please collect the command list shown and attach them to your TAC service request.

<#root>

```
show tech-support
```

```
show platform hardware fed active fwd-asic resource tcam utilization
```

```
show platform hardware fed active fwd-asic resource utilization
```

```
show platform hardware fed active fwd-asic resource rewrite utilization
```

```
show platform hardware fed active fwd-asic resource features ip-adjacency utilization
```

```
show platform software fed active ip route summary
```

```
show platform software fed active hash l3unicast
```

```
show ip route
```

```
show ip route vrf *
```

If you find an issue with this document use the **Feedback** button located at the right panel of this article. Provide as much detail as possible when you submit the request or feedback on this document, which includes information about the section, area, or issue you had with the document and what could be improved.

Related Information

Technical Support & Documentation - Cisco Systems

Understand Hardware Resources on Catalyst 9000 Switches

IP Routing Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9500 Switches) - Chapter: Configuring OSPF

IP Routing Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9500 Switches) - Chapter: Configuring EIGRP

IP Routing Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9500 Switches) - Chapter: Configuring BGP

System Management Configuration Guide, Cisco IOS XE Amsterdam 17.3.x (Catalyst 9500 Switches) - Chapter: Configuring SDM Templates

Cisco Catalyst 9200 Series Switches Data Sheet

Cisco Catalyst 9300 Series Switches Data Sheet

Cisco Catalyst 9400 Series Switches Data Sheets

Cisco Catalyst 9500 Series Switches Data Sheets

Cisco Catalyst 9600 Series Switches Data Sheet

Cisco Catalyst 9500 Architecture White Paper