



Graceful Handling of Out of Resource Situations

Out of Resource Situation is an alarm or notification indicating that the resources of the router are used extensively and the resources are reaching their threshold limits. These situations can occur due to various reasons such as high number of routes, MAC addresses, interfaces, and Access Control List (ACEs). To exemplify, if the router fails to program high number of routes in its Network Processing Unit (NPU), OOR alarm can be triggered. The OOR situations ultimately leads to traffic loss. By handling the OOR situation gracefully, you can avoid the traffic loss in the router.

NPU is an integrated circuit which has a feature set targeted at a networking application domain. The role of NPU is similar to the role of CPU (Central Processing Unit) in the computer. Integrated circuits in NPU handle data packets transmission in the routers. To enable the transmission of data packets, NPU uses several internal resources such as Forwarding Information Base (FIB), Input Logical Interface (INLIF), INLIF1, INLIF2, and Router Interface (RIF).

To ensure optimum traffic forwarding without any loss, it is crucial to monitor and ensure that the NPU tables are not out of resources. The router maintains default OOR threshold levels to alert you of the NPU resource usage.

The default values for OOR threshold levels are as follows:

- The **Red** state occurs when 95% or more of the router's hardware resources are in-use.
- The **Yellow** state occurs when 80% or more of the router's hardware resources are in-use.
- The **Green** state indicates that less than 80% of the router's hardware resources are in-use and there's favorable utilization of hardware resources.
- [Out of Resource Handling of FIB Resources, on page 1](#)
- [Out of Resource Handling of Input Logical Interface and Router Interface Resources, on page 5](#)

Out of Resource Handling of FIB Resources

The main function of the FIB within routers/line-cards is to map destination IP prefixes/labels to potential egress interfaces. In order to achieve this, the FIB maintains multiple databases or hardware tables, such as:

- Forward Equivalence Class (FEC) and the Equal Cost Multi-Path Forward Equivalence Class (ECMP FEC): This database mainly stores information related to the IPv4 and IPv6 prefixes that the router has learnt. When a traffic packet reaches the router, it performs a FIB lookup on the destination prefix of the incoming packet. This lookup results in the FEC or ECMP FEC object for that particular prefix and it

provides information about the egress interface on the router through which the packet can reach its destination.

OOB Protection Mode

When the FIB resource usage exceeds the red threshold, the system enters OOB protection mode.

In FEC OOB protection mode, the router doesn't allocate any FEC objects when it learns a new prefix, but assigns a pre-created drop FEC instead. The router drops incoming packets with destination IP addresses that were assigned the drop FEC.

In ECMP FEC OOB protection mode, the router chooses only the first path among multiple paths to forward traffic. When FEC or ECMP FEC usage drops below yellow threshold, the router reprograms the affected prefixes to normal FEC or ECMP FEC. Thereafter, the router resumes traffic flow.

Configuration Steps to Change OOB Threshold Levels

The following section shows how to change the default threshold levels for OOB handling.

```
Router# conf t
Router(config)# oor hw threshold red 90
Router(config)# oor hw threshold yellow 75
Router(config)# commit
```

Release Stale FEC Resources

To release the stale FEC resources, execute the following command:

```
Router# clear cef
```

Verification of FEC Resources

Use the `show controllers npu resources` command to verify FIB resources.

The **OOB State** in the output of the `show controllers npu resources` command changes when the router reaches an OOB situation. The **OOB State** changes from **Green** to **Yellow**, and finally to **Red** depending on the utilization of FEC or ECMP FEC resources.

```
Router# show controllers npu resources fec location 0/0/CPU0
HW Resource Information
  Name                : fec
  Asic Type            : Qumran

NPU-0
OOB Summary
  Estimated Max Entries : 126976
  Red Threshold         : 95 %
  Yellow Threshold      : 80 %
  OOB State             : Green
  Bank Info             : FEC

OFA Table Information
(May not match HW usage)
  ipnhgroup           : 43058
  ip6nhgroup          : 2
  edpl                 : 0
  limd                 : 0
```

```

punt : 19
iptunneldecap : 0
ipmcroute : 1
ip6mcroute : 0
ipnh : 0
ip6nh : 0
mplsmdtbud : 0
ipvrf : 2
ippbr : 0
redirectvrf : 0
l2protect : 0
l2bridgeport : 0

```

Current Hardware Usage

```

Name: fec
Estimated Max Entries : 126976
Total In-Use : 43082 (33 %)
OOR State : Green
Bank Info : FEC

```

```

Name: hier_0
Estimated Max Entries : 126976
Total In-Use : 43082 (33 %)
OOR State : Green
Bank Info : FEC

```

Verification of EEDB Resources

The router processes several entries and often stores these entries in the form of tables. These tables are further divided into smaller tables. These smaller tables are called as banks. The banks are often named as bank_0, bank_1, bank_2, and so on. Router segregates the entries through these banks.



Note **bank_0** will always be indicated as completely utilized and in Red OOR state on the following routers and line cards:

- NCS-5501
- NCS-5501-SE
- NCS-5502
- NCS-5502-SE
- NC55-36x100G
- NC55-18H18F
- NC55-24x100G-SE
- NC55-24H12F-SE
- NC55-36x100G-S
- NC55-6x200-DWDM-S

This complete utilization of **bank_0** must be ignored since it is reserved for internal usage regardless of the router configuration.

Use **show controllers npu resources encap** command to verify the usage of EEDB resources.

The **OOD State** in the output of the **show controllers npu resources encap** command changes when the router reaches an OOR situation. The **OOD State** changes from **Green** to **Yellow**, and finally to **Red** depending on the utilization of the EEDB resources.

```
Router# show controllers npu resources encap location 1/0/CPU0
```

```
HW Resource Information
```

```
Name                : encap
Asic Type           : Jericho
```

```
NPU-0
```

```
OOD Summary
```

```
Red Threshold      : 95 %
Yellow Threshold   : 80 %
```

```
OFA Table Information
```

```
(May not match HW usage)
```

```
ipnh                : 13
ip6nh               : 0
mplsnh              : 0
llnh                : 0
srv6nh              : 0
ipvrf               : 0
mplsmdtbud          : 0
iptunnelencap       : 0
tep                 : 0
```

```
Current Hardware Usage
```

```
Name: encap
```

```
Name: bank_0
Estimated Max Entries : 4096
Total In-Use          : 4096      (100 %)
OOD State            : Red
OOD State Change Time : 2022.Mar.15 05:33:14 UTC
Bank Info              : phase=2 extended=no
```

```
Name: bank_1
Estimated Max Entries : 4096
Total In-Use          : 4        (0 %)
OOD State              : Green
Bank Info              : phase=8 extended=no
```

```
Name: bank_2
Estimated Max Entries : 4096
Total In-Use          : 0        (0 %)
OOD State              : Green
Bank Info              : phase=0 extended=no
```

```
Name: bank_3
Estimated Max Entries : 4096
Total In-Use          : 0        (0 %)
OOD State              : Green
Bank Info              : phase=0 extended=no
```

Out of Resource Handling of Input Logical Interface and Router Interface Resources

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Out of Resource Handling of Input Logical Interface and Router Interface Resources	Release 7.7.1	<p>You can now reconfigure the threshold level for NPU resources - Input Logical Interface (INLIF) and Router Interface (RIF) by changing the predefined threshold level at which Out of Resource (OOR) situation is triggered. Graceful handling of OOR helps you to minimize traffic loss.</p> <p>You get notified via systemlogs, when the utilization of resources reaches their OOR limit. Also, you can view the utilization of resources by using the following commands:</p> <ul style="list-style-type: none"> • show controllers npu resources • show grid pool

INLIF: This resource is part of Logical Interface (LIF) and mainly stores several IP-related attributes. INLIF stores information about the following:

- Layer 2 ports
- Layer 2 interface
- IP virtual routing and forwarding
- Layer 2 bridge
- Layer 2 cross connect
- MPLS
- MPLS Multicast Distribution Tree (MDT) bud
- IP tunnel decapsulation
- Policy-based routing IPv4
- Redirect virtual routing and forwarding

RIF: This resource stores information about routing interfaces and Virtual Switching Interfaces (VSI). Routing interface consists of a port at which a router connects to the given network. RIF stores information about the following:

- Layer 3 interface
- IP virtual routing and forwarding
- IP tunnel decapsulation
- Layer 2 bridge
- Policy-based routing IPv4
- Policy-based routing IPv6
- Layer 2 bridge
- Redirect virtual routing and forwarding
- MPLS Multicast Distribution Tree (MDT) bud

Cisco IOS XR Software Release 7.7.1 introduces the graceful handling of INLIF and RIF resources under NPU. You can redefine the threshold levels of OOR to change the triggering of OOR notification.

To change OOR threshold levels for INLIF and RIF, refer section *Configuration Steps to Change OOR Threshold Levels* in this chapter.

Verification

The **OOOR State** in the output of the **show controllers npu resources** and **show grid pool** command changes when the router reaches an OOR situation. The **OOOR State** changes from **Green** to **Yellow**, and finally to **Red** depending on the utilization of INLIF and RIF resources.

Use the following commands to check the utilization of NPU resources:

- **show controllers npu resources**
- **show grid pool**

```
Router#show controllers npu resources INLIF1 location 0/3/CPU0
HW Resource Information
  Name           : INLIF1
  Asic Type      : Jericho Two

NPU-0
OOR Summary
  Estimated Max Entries : 21846
  Red Threshold        : 95 %
  Yellow Threshold     : 75 %
  OOR State            : Green
  Bank Info           : INLIF1

OFA Table Information
(May not match HW usage)
  l2port          : 12
  l2intf          : 0
  ipvrf           : 1
  l2bridge        : 0
  mplsmdbud      : 0

Current Hardware Usage
Name: INLIF1
  Estimated Max Entries : 100
  Total In-Use         : 13      (13 %)
```

```

    OOR State           : Green
    Bank Info          : INLIF1

```

```

NPU-1
  OOR Summary
    Estimated Max Entries : 21846
    Red Threshold         : 95 %
    Yellow Threshold      : 75 %
    OOR State             : Green
    Bank Info             : INLIF1

```

```

OFA Table Information
(May not match HW usage)
  l2port           : 12
  l2intf           : 0
  ipvrf            : 1
  l2bridge         : 0
  mplsmdbud       : 0

```

```

Current Hardware Usage
  Name: INLIF1
    Estimated Max Entries : 100
    Total In-Use         : 13      (13 %)
    OOR State           : Green
    Bank Info            : INLIF1

```

Router#show controllers npu resources INLIF2 location 0/3/CPU0

```

HW Resource Information
  Name           : INLIF2
  Asic Type      : Jericho Two

```

```

NPU-0
  OOR Summary
    Estimated Max Entries : 100
    Red Threshold       : 95 %
    Yellow Threshold    : 75 %
    OOR State           : Green
    Bank Info            : INLIF2

```

```

OFA Table Information
(May not match HW usage)
  mplsmdbud       : 0
  l2xc            : 0
  mplslabel       : 0
  iptunneldecap   : 0
  pbr_tt_ipv4     : 0
  redirectvrf     : 0
  mplspwepport    : 0
  srv6sid         : 0

```

```

Current Hardware Usage
  Name: INLIF2
    Estimated Max Entries : 21846
    Total In-Use         : 0      (0 %)
    OOR State             : Green
    Bank Info             : INLIF2

```

```

NPU-1
OOR Summary
  Estimated Max Entries      : 21846
  Red Threshold              : 95 %
  Yellow Threshold          : 75 %
  OOR State                  : Green
  Bank Info                  : INLIF2

OFA Table Information
(May not match HW usage)
  mplsmdbud                  : 0
  l2xc                       : 0
  mplslabel                  : 0
  iptunneldecap             : 0
  pbr_tt_ipv4               : 0
  redirectvrf               : 0
  mplspwepport              : 0
  srv6sid                   : 0

Current Hardware Usage
Name: INLIF2
  Estimated Max Entries      : 100
  Total In-Use              : 0      (0 %)
  OOR State                : Green
  Bank Info                  : INLIF2

```

You can also check the utilization of resources within a bank by using the **show grid pool** command. This command gives you detailed information about the current state of banks in each resource.



Note Yang data model support is not available for **show grid pool**.

```

Router#show grid pool 1 bank all
Tue Jul 26 11:44:24.960 UTC

Bank Ptr          : 0x308ca4bd50
Bank ID           : 0
Pool            : RIF (id 1)
Bank Start        : -1
Bank End          : -1
Max Bank Size     : 1
Max Resource Pages : 1
Available resource IDs : 1 (100.000% free)
Alarm state     : Green
Bank statistics:
  Resource IDs reserved      Success      Error      (since last clear)
  Resource IDs returned      0            0            0            0

Bank Ptr          : 0x308ca4bdb8
Bank ID           : 1
Pool              : RIF (id 1)
Bank Start        : 6
Bank End          : 8192
Max Bank Size     : 8187
Max Resource Pages : 256
Available resource IDs : 8119 (99.169% free)
Alarm state       : Green
HW Resources:
  RIF_VSI
Bank statistics:
  Resource IDs reserved      Success      Error      (since last clear)
  Resource IDs reserved      986          0            986          0

```


Resource IDs returned		918	0	918	0
Client	: vlan-fib				
Resource IDs reserved		933	0	933	0
Resource IDs returned		865	0	865	0
current usage	: 68				
Client	: ip-tunnel				
Resource IDs reserved		32	0	32	0
Resource IDs returned		32	0	32	0
current usage	: 0				
Client	: redirectvrf				
Resource IDs reserved		1	0	1	0
Resource IDs returned		1	0	1	0
current usage	: 0				
Client	: l2vpn-mgr				
Resource IDs reserved		20	0	20	0
Resource IDs returned		20	0	20	0
current usage	: 0				

System Log Alerts for INLIF and RIF OOR

When utilization of resources in NPU transition from the current OOR state to another state, the router generates system logs to alert you to free up the resources before traffic loss occurs.

```
RP/0/RP0/CPU0:Jun  2 17:54:43.264 UTC: grid_svr[194]: %L2-GRID-4-BANK_ALARM_STATE_YELLOW :
  GRID POOL: RIF (id:1), BANK 0 has state changed from Green to Yellow. LC/0/3/CPU0:
  fia_driver[231]: %PLATFORM-OFA-4-_OOR_YELLOW : NPU 0, Table mplsnh, Resource INLIF1
LC/0/3/CPU0: UTC: fia_driver[231]: %PLATFORM-OFA-4-_OOR_YELLOW : NPU 1, Table mplsnh,
Resource rif
LC/0/3/CPU0: UTC: fia_driver[231]: %PLATFORM-OFA-4-_OOR_YELLOW : NPU 0, Table mplsnh,
Resource rif
LC/0/3/CPU0:Apr 12 18:04:08.221 UTC: fia_driver[231]: %PLATFORM-OFA-5-_OOR_GREEN : NPU 1,
Table mplsnh, Resource INLIF1
```

