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 trigerred. 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
- Traffic Buffer Resource Consumption Alerts, on page 21

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.

- Egress Encapsulation Data Base (EEDB): This database stores adjacency details such as the link-local details from Address Resolution Protocol (ARP) or Neighbor Discovery (ND) Protocol, Multiprotocol Label Switching (MPLS) labels, Generic Routing Encapsulation (GRE) headers and so on.

OOR Protection Mode

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

In FEC OOR 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 OOR 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.

In EEDB OOR protection mode, the router will not program any more egress enacapsulation information. This in turn results in traffic drop. When EEDB usage drops below yellow threshold, the router resumes programing new egress encapsulation data as well as traffic flow.

Configuration Steps to Change OOR Threshold Levels

The following section shows how to change the default threshold levels for OOR 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

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Hierarchical Forward Equivalence Class (HFEC) Out of Resource (OOR) Enhancements	Release 7.5.1	With this enhancement, you can view the details of the utilization of Forwarding Information Base (FIB) hardware resources, such as hierarchical FEC and hierarchical Equal Cost Multi-Path (ECMP) FEC, in the output of the command show controllers npu resources. This feature also enables the router to display system logs on the console that alert you when FEC resources have crossed the OOR threshold levels. These logs help you to take corrective action and free up FEC resources, to minimize traffic loss. .

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

The **OOR State** in the output of the **show controllers npu resources** command changes when the router reaches an OOR situation. The **OOR 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
OOR Summary
        Estimated Max Entries       : 126976
        Red Threshold               : 95 %
        Yellow Threshold            : 80 %
        OOR 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
```

FIB has three FEC hierarchies and 3 ECMP FEC hierarchies. From Cisco IOS XR Release 7.5.1 onwards, you can view these hierarchies in the output of the commands **show controllers npu resources fec** and **show controllers npu resources ecmpfec**. The router records the **OOR State** of each of these hierarchies in the command outputs.

```
Router# show controllers npu resources fec location 1/0/CPU0
HW Resource Information
    Name                           : fec
    Asic Type                      : Jericho 2

NPU-0
OOR Summary
        Estimated Max Entries      : 576704
        Red Threshold              : 95 %
        Yellow Threshold           : 80 %
        OOR State                  : Green
        Bank Info                  : FEC


OFA Table Information
(May not match HW usage)
        ipnhgroup                  : 157800
        ip6nhgroup                 : 17377
        edpl                       : 0
        limd                       : 0
        punt                       : 18
        iptunneldecap              : 0
        ipmcroute                  : 1
        ip6mcroute                 : 0
        ipnh                       : 0
        ip6nh                      : 0
        mplsmdtbud                 : 0
        ipvrf                      : 1053
        ippbr                      : 0
        redirectvrf                : 1
        l2protect                  : 42
        l2bridgeport               : 58

Current Hardware Usage
    Name: fec
        Estimated Max Entries      : 576704
        Total In-Use               : 176350   (30 %)
        OOR State                  : Green
        Bank Info                  : FEC
```

```
        Name: hier_0
            Estimated Max Entries       : 131072
            Total In-Use                : 3207     (2 %)
            OOR State                   : Green
            Bank Info                   : H1 FEC


        Name: hier_1
            Estimated Max Entries       : 262144
            Total In-Use                : 63       (0 %)
            OOR State                   : Green
            Bank Info                   : H2 FEC


        Name: hier_2
            Estimated Max Entries       : 183488
            Total In-Use                : 173080   (94 %)
            OOR State                   : Yellow
            OOR State Change Time       : 2021.Aug.29 22:14:16 PDT
            Bank Info                   : H3 FEC
Router# show controllers npu resources ecmpfec location 0/0/CPU0
HW Resource Information
    Name                            : ecmp_fec
    Asic Type                       : Jericho 2

NPU-0
OOR Summary
        Estimated Max Entries       : 32768
        Red Threshold               : 95 %
        Yellow Threshold            : 80 %
        OOR State                   : Yellow
        OOR State Change Time       : 2021.Aug.29 23:07:53 PDT
        Bank Info                   : ECMP


OFA Table Information
(May not match HW usage)
        ipnhgroup                   : 30654
        ip6nhgroup                  : 4

Current Hardware Usage
    Name: ecmp_fec
        Estimated Max Entries       : 32768
        Total In-Use                : 30658    (93 %)
        OOR State                   : Yellow
        OOR State Change Time       : 2021.Aug.29 23:07:53 PDT
        Bank Info                   : ECMP


        Name: hier_0
            Total In-Use                : 0
            OOR State                   : Green
            Bank Info                   : H1 ECMP


        Name: hier_1
            Total In-Use                : 1
            OOR State                   : Green
            Bank Info                   : H2 ECMP


        Name: hier_2
            Total In-Use                : 30657
```

```
OOR State                   : Yellow
OOR State Change Time       : 2021.Aug.29 23:07:53 PDT
Bank Info                   : H3 ECMP
```

System Log Alerts for HFEC OOR

When utilization of any of the FEC or ECMP FEC hierarchies transition from the current OOR state to another state, the router generates system logs to alert the user. You can then take corrective action if the router is approaching OOR state to ensure that FEC resources get freed up before traffic loss takes place.

```
LC/0/0/CPU0: fia_driver[170]: %PLATFORM-OFA-1-OOR_RED : NPU 0, Table ipnhgroup, Resource
fec
LC/0/0/CPU0: fia_driver[170]: %PLATFORM-OFA-4-OOR_YELLOW : NPU 0, Table ipnhgroup, Resource
 fec
LC/0/0/CPU0: fia_driver[170]: %PLATFORM-OFA-5-OOR_GREEN : NPU 0, Table ipnhgroup, Resource
 fec
LC/0/0/CPU0: fia_driver[140]: %PLATFORM-OFA-4-OOR_YELLOW : NPU 0, Table ipnhgroup, Resource
 ecmp_fec
LC/0/0/CPU0: fib_mgr[133]: %PLATFORM-PLAT_FIB-4-OOR_PROT_STOP_WARNING : ECMP_FEC resources
 are now available. OOR protection stopped
LC/0/0/CPU0: fia_driver[140]: %PLATFORM-OFA-5-OOR_GREEN : NPU 0, Table ipnhgroup, Resource
 ecmp_fec
```

Verification of EEDB Resources

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
EEDB Resource Segregation for Encapsulation Types	Release 7.8.1	We have now segregated the Egress Encapsulation Database (EEDB) resources based on different egress encapsulation types, such as attachment circuits (AC), pseudowires, tunnels, and Address Resolution Protocol (ARP). This allows you to examine encapsulation-specific resource details using show commands. When resource usage exceeds the Out of Resource (OOR) threshold levels, the router also alerts you with system log messages. When resource usage exceeds the OOR thresholds, you can avoid traffic loss by taking corrective action to free up resources, such as reducing the scale of interfaces with the related encapsulation. The **show controllers npu resources** command is now modified to include the following optional keywords: • **encapAC** • **encapPWE** • **encaptunnels** • **encapARP** This feature is supported on Cisco 5700 Series Routers and on routers that have the NC57 line cards installed and operating in either native or compatiblity mode

Feature Name	Release Information	Feature Description
Verification of EEDB OOR	Release 7.6.1	If the Egress Encapsulation Database (EEDB) resources exceed the Out of Resource (OOR) threshold levels, the router generates system logs that indicate the resource utilization of EEDB hardware. With these details, you can take corrective action by freeing up EEDB resources, thereby minimizing traffic loss. This feature also updates the show output of the following command to display utilization of EEDB hardware resources: show controllers npu 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 **OOR State** in the output of the **show controllers npu resources encap** command changes when the router reaches an OOR situation. The **OOR 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
 OOR 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 %)
      OOR State                   : Red
      OOR 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 %)
      OOR State                   : Green
      Bank Info                   : phase=8 extended=no


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


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

From Cisco IOS XR Release 7.6.1 onwards, the output of the command **show controllers npu resources encap** has been enhanced to provide more fields such as **Current Hardware Usage** and **Cluster Bank Pair** details.

```
Router# show controllers npu resources encap location 1/0/CPU0
 HW Resource Information
  Name                            : encap
  Asic Type                       : Jericho Two
```

```
NPU-0
OOR 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 J2 Hardware Usage

   Cluster Bank Pair: EEDB_S2_XL
 Max-Entries                      : 163840
 Total In-use                     : 0

 Logical_phase                    : 1 (Encap_Rif)
  Cluster Bank
     Bank Size               : S2
     Estimated Max Entries   : 163840
     OOR State               : Green
     Total In-Use            : 0

     EEDB Bank:
    Estimated Max Entries  : 8192
    OOR State              : Green
    Total In-Use           : 0

 Logical_phase                    : 6 (Tunnel4)
  Cluster Bank
     Bank Size               : XL
     Estimated Max Entries   : 122880
     OOR State               : Green
     Total In-Use            : 0

     EEDB Bank:
    Estimated Max Entries  : 61440
    OOR State              : Green
    Total In-Use           : 0


   Cluster Bank Pair: EEDB_L1_S1
 Max-Entries                      : 98304
 Total In-use                     : 26

 Logical_phase                    : 2 (Encap_NativeArp)
  Cluster Bank
     Bank Size               : L1
     Estimated Max Entries   : 81920
     OOR State               : Green
     Total In-Use            : 0

     EEDB Bank:
    Estimated Max Entries  : 30720
    OOR State              : Green
```

```
    Total In-Use           : 0

  Logical_phase                    : 8 (Encap_Ac)
   Cluster Bank
      Bank Size                : S1
      Estimated Max Entries    : 98304
      OOR State                : Green
      Total In-Use             : 26

      EEDB Bank:
     Estimated Max Entries  : 8192
     OOR State              : Green
     Total In-Use           : 0


    Cluster Bank Pair: EEDB_M1_M2
  Max-Entries                      : 131072
  Total In-use                     : 62

  Logical_phase                    : 3 (Encap_NativeAc or Tunnel1)
   Cluster Bank
      Bank Size                : M1
      Estimated Max Entries    : 122880
      OOR State                : Green
      Total In-Use             : 0

      EEDB Bank:
     Estimated Max Entries  : 16384
     OOR State              : Green
     Total In-Use           : 0
```

EEDB Resource Segregation for Encapsulation Types

Starting from Cisco IOS XR Release 7.8.1, we have segregated EEDB resources based on different egress encapsulation types. Cisco IOS XR software now provides encapsulation specific EEDB resource usage for the following egress encapsulation types:

- Attachment Circuits (AC)
- Pseudowires
- ARP packets
- Tunnels

To view the EEDB resource details for AC, use the command **show controllers npu resources encapAC**.

```
Router# show controllers npu resources encapAC location 0/RP0/CPU0
HW Resource Information
    Name                           : encap_AC
    Asic Type                      : Jericho Two

NPU-0
OOR Summary
        Red Threshold              : 95 %
        Yellow Threshold           : 80 %


OFA Table Information
(May not match HW usage)
        ip6nh                      : 0
        ipmctxintf                 : 0
        l2intf                     : 0
```

```
        l2port                      : 0

Current J2 Hardware Usage

   Cluster Bank Pair: EEDB_S1_L1
    Max-Entries                     : 98304
    Total In-use                    : 18

    Logical_phase                   : 8 (Encap_Ac)
        Cluster Bank
           Bank Size                : S1
           Estimated Max Entries    : 98304
           OOR State                : Green
           Total In-Use             : 18

           EEDB Bank:
             Estimated Max Entries  : 8192
             OOR State              : Green
             Total In-Use           : 0

    Logical_phase                   : 2 (Encap_NativeArp)
        Cluster Bank
           Bank Size                : L1
           Estimated Max Entries    : 81920
           OOR State                : Green
           Total In-Use             : 0

           EEDB Bank:
             Estimated Max Entries  : 30720
             OOR State              : Green
             Total In-Use           : 0
```

To view the EEDB resource details for pseudowires, use the command **show controllers npu resources encapPWE**.

```
Router# show controllers npu resources encapPWE location 0/RP0/CPU0
HW Resource Information
    Name                            : encap_PWE
    Asic Type                       : Jericho Two

NPU-0
OOR Summary
        Red Threshold               : 95 %
        Yellow Threshold            : 80 %


OFA Table Information
(May not match HW usage)
        mplspweport                 : 0
        l2bridgeolist               : 0

Current J2 Hardware Usage

   Cluster Bank Pair: EEDB_L1_S1
    Max-Entries                     : 98304
    Total In-use                    : 18

    Logical_phase                   : 2 (Encap_NativeArp)
        Cluster Bank
           Bank Size                : L1
           Estimated Max Entries    : 81920
           OOR State                : Green
           Total In-Use             : 0
```

```
            EEDB Bank:
              Estimated Max Entries  : 30720
              OOR State              : Green
              Total In-Use           : 0

    Logical_phase                    : 8 (Encap_Ac)
        Cluster Bank
           Bank Size                 : S1
           Estimated Max Entries     : 98304
           OOR State                 : Green
           Total In-Use              : 18

           EEDB Bank:
             Estimated Max Entries  : 8192
             OOR State              : Green
             Total In-Use           : 0
```

To view the EEDB resource details for tunnels, use the command **show controllers npu resources encaptunnels**.

```
Router# show controllers npu resources encaptunnels location 0/RP0/CPU0
HW Resource Information
    Name                             : encap_tunnels
    Asic Type                        : Jericho Two

NPU-0
OOR Summary
        Red Threshold                : 95 %
        Yellow Threshold             : 80 %


OFA Table Information
(May not match HW usage)
        mplsnh                       : 0
        mplspweport                  : 0
        iptunnelencap                : 0
        limd                         : 0
        ipmcmdtencap                 : 0
        srv6nh                       : 0

Current J2 Hardware Usage

   Cluster Bank Pair: EEDB_S2_XL
    Max-Entries                      : 163840
    Total In-use                     : 0

    Logical_phase                    : 1 (Encap_Rif)
        Cluster Bank
           Bank Size                 : S2
           Estimated Max Entries     : 163840
           OOR State                 : Green
           Total In-Use              : 0

           EEDB Bank:
             Estimated Max Entries  : 8192
             OOR State              : Green
             Total In-Use           : 0

    Logical_phase                    : 6 (Tunnel4)
        Cluster Bank
           Bank Size                 : XL
           Estimated Max Entries     : 122880
           OOR State                 : Green
           Total In-Use              : 0
```

```
        EEDB Bank:
          Estimated Max Entries  : 63488
          OOR State              : Green
          Total In-Use           : 0


Cluster Bank Pair: EEDB_M1_M2
 Max-Entries                     : 131072
 Total In-use                    : 58

 Logical_phase                   : 3 (Encap_NativeAc or Tunnel1)
     Cluster Bank
        Bank Size                : M1
        Estimated Max Entries    : 122880
        OOR State                : Green
        Total In-Use             : 0

        EEDB Bank:
          Estimated Max Entries  : 16384
          OOR State              : Green
          Total In-Use           : 0

 Logical_phase                   : 7 (Encap_Arp)
     Cluster Bank
        Bank Size                : M2
        Estimated Max Entries    : 131072
        OOR State                : Green
        Total In-Use             : 58

        EEDB Bank:
          Estimated Max Entries  : 16384
          OOR State              : Green
          Total In-Use           : 0


Cluster Bank Pair: EEDB_L2_M3
 Max-Entries                     : 65536
 Total In-use                    : 2

 Logical_phase                   : 4 (Encap_Tunnel2)
     Cluster Bank
        Bank Size                : L2
        Estimated Max Entries    : 57344
        OOR State                : Green
        Total In-Use             : 0

        EEDB Bank:
          Estimated Max Entries  : 30720
          OOR State              : Green
          Total In-Use           : 0

 Logical_phase                   : 5 (Tunnel3)
     Cluster Bank
        Bank Size                : M3
        Estimated Max Entries    : 65536
        OOR State                : Green
        Total In-Use             : 2

        EEDB Bank:
          Estimated Max Entries  : 16384
          OOR State              : Green
```

```
            Total In-Use           : 0
```

To view the EEDB resource details for ARP, use the command **show controllers npu resources encapARP**.

```
Router# show controllers npu resources encapARP location 0/RP0/CPU0
HW Resource Information
    Name                           : encap_ARP
    Asic Type                      : Jericho Two

NPU-0
OOR Summary
        Red Threshold              : 95 %
        Yellow Threshold           : 80 %


OFA Table Information
(May not match HW usage)
        ipvrf                      : 0
        redirectvrf                : 0
        ipnh                       : 0
        ip6nh                      : 0
        mplsmdtbud                 : 0
        llnh                       : 0

Current J2 Hardware Usage

   Cluster Bank Pair: EEDB_M2_M1
    Max-Entries                    : 131072
    Total In-use                   : 58

    Logical_phase                  : 7 (Encap_Arp)
        Cluster Bank
           Bank Size               : M2
           Estimated Max Entries   : 131072
           OOR State               : Green
           Total In-Use            : 58

           EEDB Bank:
             Estimated Max Entries : 16384
             OOR State             : Green
             Total In-Use          : 0

    Logical_phase                  : 3 (Encap_NativeAc or Tunnel1)
        Cluster Bank
           Bank Size               : M1
           Estimated Max Entries   : 122880
           OOR State               : Green
           Total In-Use            : 0

           EEDB Bank:
             Estimated Max Entries : 16384
             OOR State             : Green
             Total In-Use          : 0
```

System Log Alerts for EEDB OOR

When utilization of EEDB resources transition from the current OOR state to another state, the router generates system logs to alert the user. You can then take corrective action if the router is approaching OOR state to ensure that EEDB resources get freed up before traffic loss takes place.

The router logs the following log messages to inform you of the OOR state of the encap resource when it exceeds thresholds. The syslog also indicates the Network Processing Unit (NPU) of the affected encap resource, which is NPU0, in this case.

```
LC/0/0/CPU0: fia_driver[164]: %PLATFORM-OFA-4-OOR_YELLOW : NPU 0, Table mplsnh, Resource
encap
LC/0/0/CPU0: fia_driver[164]: %PLATFORM-OFA-1-OOR_RED : NPU 0, Table mplsnh, Resource encap
```

Starting from Cisco IOS XR Release 7.8.1, the router logs the following GRID alarm log message and NPU OOR alarm log message when the encapsulation specific EEDB resource exceeds thresholds.

```
LC/0/0/CPU0: grid_svr[450]: %L2-GRID-4-BANK_ALARM_STATE_YELLOW : GRID POOL: GLIF (id:2),
BANK 0 has state changed from Green to Yellow.
LC/0/0/CPU0: fia_driver[228]: %PLATFORM-OFA-4-OOR_YELLOW : NPU 0, Table ipnh, Resource
encap_ARP
```

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

Table 3: 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	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. 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: • 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 **OOR State** in the output of the **show controllers npu resources** and **show grid pool**command changes when the router reaches an OOR situation. The **OOR 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
        mplsmdtbud                  : 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
        mplsmdtbud                  : 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)
        mplsmdtbud                  : 0
        l2xc                        : 0
        mplslabel                   : 0
```

```
        iptunneldecap               : 0
        pbr_tt_ipv4                 : 0
        redirectvrf                 : 0
        mplspweport                 : 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)
        mplsmdtbud                  : 0
        l2xc                        : 0
        mplslabel                   : 0
        iptunneldecap               : 0
        pbr_tt_ipv4                 : 0
        redirectvrf                 : 0
        mplspweport                 : 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:                          Success      Error         (since last clear)
  Resource IDs reserved                         0          0                 0          0
```

```
  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:                          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
```

Traffic Buffer Resource Consumption Alerts

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Traffic Buffer Resource Consumption Alerts	Release 7.7.1	You can now configure threshold values for available traffic buffer resources and get timely syslog alerts on the router console when available resources go below the configured threshold values. These notifications enable you to free up resources or reroute traffic before the router drops traffic packets due to resource exhaustion. In earlier releases, the router dropped traffic without warning when traffic buffer resources got exhausted. This feature is supported on Cisco 5700 Series Routers and routers that have the NC57 line cards installed and operating in either native or compatibility mode. This feature introduces the **hw-module profile qos free-buffer-int-threshold** command

The main functionality of the Network Processor Unit (NPU) of your router is to forward incoming traffic packets and manage traffic. The TM module of the NPU takes care of traffic management. TM enforces Quality of Service (QoS) and buffer management of the data traffic that traverses the router. TM has three kinds of buffer resources:

- Static Random Access Memory (SRAM) buffer: This on-chip memory stores packet data that the NPU has to access with minimum latency.
- SRAM Packet Descriptor Buffer (PDB): This on-chip memory stores descriptor information for data traffic packets such as the packet length, memory location of the packet, etc.
- Dynamic Random Access Memory (DRAM) Bundle Descriptor Buffer (BDB): When SRAM resources become scarce, TM moves packets from on-chip SRAM to this off-chip memory and stores the packets as DRAM bundles. DRAM is slower compared to SRAM.

The congestion module of the TM tracks the buffer usage. IOS XR operating system displays a system log message on the console when free buffers are lower than the configured threshold value. This log message enables you to take necessary action to free up buffer resources. If you take no action, then the router drops traffic when all the free buffers get exhausted.

Configure Threshold Values for free TM Buffer Resources

You can configure the threshold values for free TM buffer resources using the command **hw-module profile qos free-buffer-int-threshold** *set-value clear-value* as shown in the following configuration example:

```
Router# conf t
Router(config)# hw-module profile qos free-buffer-int-threshold 50 75
Router(config)# commit
Router# reload location all
```

Note

- After configuring the thresholds, reload all locations to apply the configured values.
- To disable this feature, use the command **no hw-module profile qos free-buffer-int-threshold** *set-value clear-value* and reload all locations.

In the preceding configuration example, when free buffers go below the *set-value*, that is, less than 50%, the TM congestion manager sets an interrupt. An interrupt is a signal sent from hardware or software to the operating system (OS) so that the OS takes immediate action. Here, the interrupt triggers the IOS XR OS to display system log messages on the console which alert you that free buffer resources are low. The TM congestion manager clears the interrupt only after the free buffers have exceeded the *clear-value* that is 75%. The TM congestion manager doesn't generate any more interrupts until the free buffers exceed the *clear-value* and then again reduces below the *set-value*.

Running Configuration

```
Router# show running-config hw-module profile qos
hw-module profile qos free-buffer-int-threshold 50 75
```

Guidelines for Configuring TM Buffer Thresholds

- Configure the *set-value* and *clear-value* within the range of 0-100, as they are percentage values of free buffer resources. The TM congestion module calculates the absolute values automatically.
- Configure the *clear-value* higher than the *set-value*.
- The configured threshold values apply to all three TM buffer resources.
- If you configure only *set-value*, the *clear-value* threshold defaults to 100.
- This feature is not enabled by default.

System Logs for TM Buffer Resource Congestion

The following code block shows the system log message that the router displays on the router console when free buffers go below the *set-value* for the SRAM buffer.

```
LC/0/7/CPU0: fia_driver[219]: %FABRIC-FIA_DRVR_ASIC-4-JER2_CGM_GLBL_RES_EXHAUST : [3186] :
 CGM_0 Interrupt for Global SRAM BUFFER observed on unit 0.
```

In the above system log message:

- LC0/7/CPU0 indicates that the message is from line card in slot 7.
- CGM_0 indicates that NPU core 0 has resource congestion.

- SRAM BUFFER is the resource that is congested
- Unit 0 indicates NPU 0.

The router displays similar system log messages for SRAM PDB and DRAM BDB as well, when free buffers go below the *set-value*.

```
LC/0/7/CPU0: fia_driver[219]: %FABRIC-FIA_DRVR_ASIC-4-JER2_CGM_GLBL_RES_EXHAUST : [3186] :
 CGM_0 Interrupt for Global SRAM PACKET DESCRIPTOR observed on unit 0.
LC/0/7/CPU0: fia_driver[219]: %FABRIC-FIA_DRVR_ASIC-4-JER2_CGM_GLBL_RES_EXHAUST : [3186] :
 CGM_0 Interrupt for Global DRAM BUNDLE DESCRIPTOR observed on unit 0.
```

Verification of TM Buffer Congestion

You can view the interrupts using the show command **show asic-errors fia all** .

```
Router# show asic-errors fia all all location 0/0/CPU0
FullQ, S9700-53DX-J, 0/0/CPU0, fia[0]
Name            : CGM_0.Congestion_Interrupt_Register.GlblSramBuffersFcInt
Leaf ID         : 0x3602803e
Error count     : 399
Last clearing   : Wed Oct  6 05:23:21 2021
Last N errors   : 50
---------------------------------------------------------------
First N errors.
@Time, Error-Data
----------------------------------------
Oct  6 05:23:21.517327
Oct  6 05:23:21.757469
Oct  6 05:23:21.759296
Oct  6 05:23:21.759752
Oct  6 05:23:21.760754
Oct  6 05:23:21.760924
Oct  6 05:23:21.761135
Oct  6 05:23:21.762766
Oct  6 05:23:21.763298
Oct  6 05:23:21.763734
Oct  6 05:23:21.763869
```