# Capture for-US Traffic with the 8000 Series Router

## Contents



## Introduction

This document describes how to capture for-us traffic in the Cisco 8000 Series router.

## Prerequisites

### Requirements

Familiarity with Cisco 8000 Series routers and Cisco IOS® XR software.

### Components Used

The information in this document is based Cisco 8000 Series routers and is not restricted to specific software and hardware version.

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

During troubleshooting activities, there are cases where you need to verify the traffic that is being switched to the Central Processing Unit (CPU ) for further processing or handling.

This article is intended to explain how this traffic can be captured in the Cisco 8000 Series router.

## Procedure

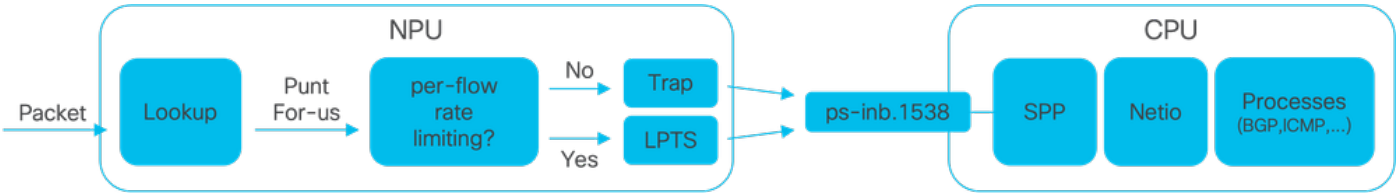

*Image1 - Cisco 8000 Series router simplified NPU and CPU diagram.*

When a packet is received In the Cisco 8000 router, a lookup is done by the Network Processing Unit (the NPU) which results in a forwarding decision.

There can be a case where the decision is to punt the packet, meaning to switch the packet to the CPU for further processing or handling.

The NPU lookup also determines if per-flow-rate-limiting is required while switching the packet to the CPU.

- If per-flow-rate-limiting is required, then the packet is switched to the CPU via the Local Packet Transport Service (LPTS), for example, a routing protocol packet.
- If per-flow-rate-limiting is not required, then a trap is generated and the packet is switched to the CPU, for example, a packet with Time-to-Live (TTL) expired.

The packets, if not rate-limited, are switched to the CPU via a dedicated internal VLAN with id 1538.

You can verify both the LPTS table and the Traps table entries using the **show lpts pifib hardware entry brief** and the **show controllers npu stats traps-all** commands.

The **show lpts pifib hardware entry brief** command displays the LPTS table entries.

Here, the output is limited to entries associated with the Border Gateway Protocol (BGP).

```
RP/0/RP0/CPU0:8202#show lpts pifib hardware entry brief location 0/rp0/cpu0 | include "Type|BGP"

Type DestIP              SrcIP               Interface          vrf   L4   LPort/Type   RPort  npu  F
IPv4 10.4.11.2           10.4.11.3           any                0     6    Port:20656   179    0    B
IPv4 10.4.11.2           10.4.11.3           any                0     6    Port:179     0      0    B
IPv4 any                 any                 any                0     6    Port:any     179    0    B
IPv4 any                 any                 any                0     6    Port:179     0      0    B
IPv6 any                 any                 any                0     6    Port:any     179    0    B
IPv6 any                 any                 any                0     6    Port:179     0      0    B
RP/0/RP0/CPU0:8202#
```

The **show controllers npu stats traps-all** command lists all traps entries and associated counters.

Here, the output is limited to entries with packet matches excluding all entries showing zero in the Packets Accepted and the Packets Dropped columns.

Note that all traps are rate-limited.

```
show controllers npu stats traps-all instance 0 location 0/rp0/cpu0 | exclude "0 0"

RP/0/RP0/CPU0:8202#show controllers npu stats traps-all instance 0 location 0/rp0/cpu0 | exclude "0

Traps marked (D*) are punted (post policing) to the local CPU internal VLAN 1586 for debugging
They can be read using "show captured packets traps" CLI
Traps marked (D) are dropped in the NPU
Traps punted to internal VLAN 1538 are processed by the process "spp" on the "Punt Dest" CPU
They can also be read using "show captured packets traps" CLI
"Configured Rate" is the rate configured by user (or default setting) in pps at the LC level
```

```
"Hardware Rate" is the actual rate in effect after hardware adjustments
Policer Level:
        NPU: Trap meter is setup per NPU in packets per second
        IFG: Trap meter is setup at every IFG in bits per second
        The per IFG meter is converted from the user configured/default rate (pps)
        based on the "Avg-Pkt Size" into bps.
        Due to hardware adjustments, the "Configured Rate" and
        "Hardware Rate" differ in values.


NOTE:The displayed stats are NOT real-time and are updated every 30 SECONDS from the hardware.

Trap Type                                       NPU  Trap  Punt       Punt  Punt  Punt Configured Hardwar
                                                ID   ID    Dest       VoQ   VLAN  TC   Rate(pps)  Rate(pp
================================================================================================================
ARP                                             0    3     RPLC_CPU   271   1538  7    542        533
NOT_MY_MAC(D*)                                  0    4     RPLC_CPU   264   1586  0    67         150
DHCPV4_SERVER                                   0    8     RPLC_CPU   265   1538  1    542        523
LLDP                                            0    26    RPLC_CPU   270   1538  6    4000       3862
ONLINE_DIAG                                     0    31    RPLC_CPU   271   1538  7    4000       3922
V4_MCAST_DISABLED(D*)                           0    69    RPLC_CPU   269   1586  5    67         150
V6_MCAST_DISABLED(D*)                           0    80    RPLC_CPU   264   1586  0    67         150
L3_IP_MULTICAST_NOT_FOUND(D*)                   0    125   RPLC_CPU   264   1586  0    67         150
RP/0/RP0/CPU0:8202#
```

The shell utility **spp_platform_pcap** can be used to capture packets that cross this dedicated internal VLAN between the NPU and the CPU. This same utility also allows to capture the traffic sent or received through the router management interface.

The **spp_platform_pcap** shell utility is executed from within the shell and provides multiple usage options. To access or login to the shell, execute the **run** command. To logout from shell, type **exit**.

```
RP/0/RP0/CPU0:8202#run
[node0_RP0_CPU0:~]$spp_platform_pcap -h
Usage: spp_platform_pcap options
Use Ctrl-C to stop anytime
  -h  --help            Display this usage information.
  -D  --Drop            capture Drops in SPP.
  -i  --interface       Interface-name
                        Available from the output of
                        "show ipv4 interface brief"
  -Q  --direction       direction of the packet
                        Options: IN | OUT |
                        Mandatory option
                        (when not using the -d option)
  -s  --source          Originator of the packet.
                        Options: ANY | CPU | NPU | NSR | MGMT | PTP | LC_PKTIO | LC_REDIR
  -d  --destination     destination of the packet
                        Options: ANY | CPU | NPU | MGMT | PTP | LC_PKTIO | LC_REDIR |
  -l  --l4protocol      IANA-L4-protocol-number
                         (use with Address family (-a)
                        Interface (-i) and direction (-Q)
                        Options: min:0 Max:255
  -a  --addressFamily   address Family used with l4protocol (-l)
                        Interface (-i) and direction (-Q)
                        Options: ipv4 | ipv6 |
  -x  --srcIp           Src-IP (v4 or v6)
                        Used with -a, -i and -Q only
```

```
  -X  --dstIp           Dst-IP (v4 or v6)
                        Used with -a, -i and -Q only
  -y  --srcPort         Src-Port
                        Used with -a, -l, -i and -Q only
                        Options: min:0 Max:65535
  -Y  --dstPort         Dst-Port
                        Used with -a, -l, -i and -Q only
                        Options: min:0 Max:65535
  -P  --l2Packet        Based on L2 packet name/etype
                        Interface (-i) and direction (-Q) needed
                        Use for non-L3 packets
                        Options:ether-type (in hex format)
                        ARP | ISIS | LACP | SYNCE | PTP | LLDP | CDP |
  -w  --wait            Wait time(in seconds)
                        Use Ctrl-C to abort
  -c  --count           Count of packets to collect
                        min:1; Max:1024
  -t  --trapNameOrId    Trap-name(in quotes) or number(in decimal)
                        (direction "in" is a MUST).
                        Refer to "show controllers npu stats traps-all instance all location <LC|RP>
                        Note: Trap names with (D*) in the display are not punted to SPP.
                        They are punted to ps-inb.1586
  -S  --puntSource      Punt-sources
                        Options: LPTS_FORWARDING | INGRESS_TRAP | EGRESS_TRAP | INBOUND_MIRROR |
                        NPUH |
  -p  --pcap            capture packets in pcap file.
  -v  --verbose         Print the filter offsets.
[node0_RP0_CPU0:~]$
```

Note the capture direction option, **-Q**, where the value **IN** means that it captures the punted packets (the packets received by the CPU). The value **OUT** means that it captures the injected packets (the packets sent by the CPU). The option **-p** allows to capture packets in a pcap file.

Please consider that, by default, the **spp_platform_pcap** capture:

- Runs for 60 seconds.
- Captures a maximum of 100 packets.
- Truncates all captured packets to 214 Bytes.

For example, to start an unfiltered capture of all traffic received by the CPU, type the command **spp_platform_pcap -Q IN -p**:

```
[node0_RP0_CPU0:~]$spp_platform_pcap -Q IN -p
All trace-enabled SPP nodes will be traced.
Node "socket/rx" set for trace filtering. Index: 1
Wait time is 60 seconds. Use Ctrl-C to stop
Collecting upto 100 packets (within 60 seconds)
^CSignal handling initiated <<<<<<<< Here: 'Ctrl-C' was used to stop the capture.
Tracing stopped with 10 outstanding...
Wrote 90 traces to /tmp/spp_bin_pcap
All trace-enabled SPP nodes will be traced.
pcap: Captured pcap file for packets saved at "/tmp/spp_pcap_capture_0_RP0_CPU0.pcap"

[node0_RP0_CPU0:~]$
```

When the capture ends, the resulting file is made available on the local disk.

Copy the file from the router to your local computer and verify its contents using your preferred packet decoder application.

```
[node0_RP0_CPU0:~]$ls -la /tmp
total 44
<snip>
-rw-r--r--.  1 root  root  8516 Aug  7 06:58 spp_pcap_capture_0_RP0_CPU0.pcap
<snip>
[node0_RP0_CPU0:~]$
[node0_RP0_CPU0:~]$cp /tmp/spp_pcap_capture_0_RP0_CPU0.pcap /harddisk:/
[node0_RP0_CPU0:~]$exit
logout

RP/0/RP0/CPU0:8202#dir harddisk: | include spp_pcap

16 -rw-r--r--. 1 8516 Aug 8 07:01 spp_pcap_capture_0_RP0_CPU0.pcap
RP/0/RP0/CPU0:8202#
```

It is possible to be more specific with regard to the intent of your capture. For example, you can leverage the utility filter capabilities to capture the for-us traffic related with a specific router interface, or an IP address, or a certain protocol.

As an example, using this command, you can capture the BGP traffic from a specific peer on a specific interface:

```
spp_platform_pcap -Q IN -a ipv4 -l 6 -i HundredGigE0/0/0/1 -x 10.100.0.1 -Y 179 -p
```

You can also use **spp_platform_pcap** to capture the traffic sent or received through the router management interface.

As and example, using this command, you can capture the traffic received from the management interface.

```
spp_platform_pcap -Q IN -p -i MgmtEth0/RP0/CPU0/0
```

All previous examples were executed on a standalone Cisco 8000 series router. If working with a distributed Cisco 8000 series router, consider in which node, route-processor, or linecard, you wish the capture to be executed.

It can be the case that the particular traffic you are interested in is handled by a certain linecard CPU. Both the **show controllers npu stats traps-all** and the **show lpts pifib hardware entry brief** can help identify the punt destination.

<#root>

```
RP/0/RP0/CPU0:8808#show controllers npu stats traps-all instance 0 location 0/0/cpu0 | include "Type|Ac
```

```
Trap Type                                   NPU  Trap
```

**Punt**

```
       Punt  Punt  Punt Configured Hardware   Policer Avg-Pkt Packets              Packets
                                              ID   ID
```

**Dest**

```
       VoQ   VLAN  TC   Rate(pps)  Rate(pps)  Level   Size    Accepted             Dropped
=============================================================================================================
ARP                                         0    10   LC_CPU      239   1538  7    542        531
ISIS/L3                                     0    129  BOTH_RP-CPU 239   1538  7    10000      9812
RP/0/RP0/CPU0:8808#

RP/0/RP0/CPU0:8808#show lpts pifib hardware entry brief location 0/0/cpu0 | include "Type|--|Fragment|O

-------------------------------------------------------------------------
Type DestIP               SrcIP                Interface          vrf   L4   LPort/Type   RPort  npu  F
```

**DestNode**

```
     PuntPrio       Accept Drop
---- ------------------- -------------------- --------------     ----- ---  ------------ ------ ---- -
IPv4 any                 any                  any                0     0    any          0      0    F
IPv4 any                 any                  any                0     0    any          0      0    F
IPv4 any                 any                  any                0     0    any          0      1    F
IPv4 any                 any                  any                0     0    any          0      1    F
IPv4 any                 any                  any                0     0    any          0      2    F
IPv4 any                 any                  any                0     0    any          0      2    F
IPv4 any                 any                  any                0     89   any          0      0    O
IPv4 any                 any                  any                0     89   any          0      0    O
IPv4 any                 any                  any                0     89   any          0      1    O
IPv4 any                 any                  any                0     89   any          0      2    O
IPv4 any                 any                  any                0     89   any          0      0    O
IPv4 any                 any                  any                0     89   any          0      0    O
IPv4 any                 any                  any                0     89   any          0      1    O
IPv4 any                 any                  any                0     89   any          0      2    O
IPv6 any                 any                  any                0     0    any          0      0    F
IPv6 any                 any                  any                0     0    any          0      1    F
IPv6 any                 any                  any                0     0    any          0      2    F
IPv6 any                 any                  any                0     89   any          0      0    O
IPv6 any                 any                  any                0     89   any          0      1    O
IPv6 any                 any                  any                0     89   any          0      2    O
IPv6 any                 any                  any                0     89   any          0      0    O
IPv6 any                 any                  any                0     89   any          0      1    O
IPv6 any                 any                  any                0     89   any          0      2    O
RP/0/RP0/CPU0:8808#
```

Once identified, attach to the specific linecard, and from there, execute the **spp_platform_pcap** utility as shown before.

```
attach location 0/0/cpu0
spp_platform_pcap -Q IN -p
! --- execute 'Ctrl-C' to stop the capture
```

## Related Information

*Cisco Techcnical Assistance Center (TAC) Video*

*Cisco 8000 Series - Capture for-us traffic, Video*