



Cisco IOS XRv 9000 Data Plane Specific Features

- [Early Fast Discard, on page 1](#)
- [CEF load balancing with GTP, on page 2](#)
- [Data Plane Management, on page 3](#)

Early Fast Discard

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Early Fast Discard	Release 24.1.1	Starting from Cisco IOS XR Release 24.1.1, the early fast discard feature is deprecated and will not be supported in future releases. We recommend not to use this feature starting from Cisco IOS XR Release 24.1.1.

The early fast discard (EFD) is a congestion protection feature to handle cases when the incoming traffic exceeds the router capacity. While the normal congestion control is handled by the traffic manager (TM) and attached QoS policies, the EFD feature is activated in extreme cases. The EFD feature filters high priority traffic (such as keepalives, control, and BFD) and discards the rest at the very early stage of datapath processing, thus maintaining the important flow of control packets.

To activate the EFD feature, use the **early-fast-discard** command in configuration mode. The traffic to be discarded is defined by setting IP precedence, MPLS exp, and VLAN cos values.

Configuring Early Fast Discard

Use Case

Early Fast Discard is configured on the data plane to manage incoming traffic flow by defining the traffic to be discarded based on these conditions:

- IP precedence=4
- MPLS exp=3
- VLAN cos=5

The default value is 6 and ge (greater than or equal to).

Configuration

```
Router# configure
Router(config)# hw-module early-fast-discard
Router(config-early-fast-discard)# ip-prec 4 ip-op [lt | ge]
Router(config-early-fast-discard)# mpls-exp 3 mpls-op [lt | ge]
Router(config-early-fast-discard)# vlan-cos 5 vlan-op [lt | ge]
Router(config-early-fast-discard)# exit
```



Note Use the **no hw-module early-fast-discard** command to deactivate EFD.

Running Configuration

```
RP/0/RP0/CPU0:ios#show run hw-module early-fast-discard
Thu Jul 16 15:51:34.672 UTC
hw-module early-fast-discard
 ip-prec 4 ip-op ge
 mpls-exp 3 mpls-op ge
 vlan-cos 5 vlan-op lt
!
```

Related Topics

- [Early Fast Discard, on page 1](#)

CEF load balancing with GTP

This section describes per-flow load balancing.

Per-Flow Load Balancing

Load balancing describes the functionality in a router that distributes packets across multiple links based on Layer 3 (network layer) and Layer 4 (transport layer) routing information. If the router discovers multiple paths to a destination, the routing table is updated with multiple entries for that destination.

Per-flow load balancing performs the following functions:

- Incoming data traffic is evenly distributed over multiple equal-cost connections.
- Layer 3 (network layer) load balancing decisions are taken on IPv4, IPv6, and MPLS flows which are supported for the 5-tuple hash algorithm.
- The same hash algorithm (3-tuple or 5-tuple) is used for load balancing over multiple equal-cost Layer 3 (network layer) paths. The Layer 3 (network layer) path is on a physical interface or on a bundle interface.

- The cef load-balancing fields L3 global command allows you to select the 3-tuple hash algorithm.
- By default, 5-tuple hash algorithm is used for load balancing. If you use the cef load-balancing fields L3 global command, 3-tuple hash algorithm is enabled.

Layer 3 (L3) Routing Information

The 3-tuple load-balance hash calculation contains the following Layer 3 (Network Layer) inputs:

- Source IP address
- Destination IP address
- Router ID

Layer 4 (L4) Routing Information

The 5-tuple load-balance hash calculation contains 3-tuple inputs and these additional following Layer 4 (Transport Layer) inputs:

- Source IP address
- Destination IP address

CEF load balancing for GTP is always enabled by default on XRv9K platform and cannot be disabled. Therefore, the no cef loadbalancing fields l4 gtp command does not disable the inclusion of GTP fields in hash calculation.

XRv9k can detect GTP headers even if it's behind a GRE header + MPLS Labels and can perform load balancing based on GTP fields.

Data Plane Management

Cisco IOS XRv 9000 Router's data plane starts automatically after the router is booted. In case the control plane does not establish or loses communication with the data plane, then the control plane automatically restarts the data plane.

The data plane can be started, shut-down, and reloaded manually through the admin console for maintenance and troubleshooting purposes. This table lists the commands that are required to manage the data plane.

Table 2: Data Plane Management Commands

Task	Use this command
Check the status of data plane.	<pre>show sdr sysadmin-vm:0_RP0# show sdr Thu May 7 18:38:38.996 UTC sdr default-sdr location 0/RP0/VM1 sdr-id 2 IP Address of VM 192.0.0.4 MAC address of VM E2:3A:DD:0A:8C:06 VM State RUNNING start-time 2015-05-07T17:54:39.457822+00:00 Last Reload Reason FIRST_BOOT Reboot Count 1 location 0/RP0/VM2 sdr-id 2 IP Address of VM 192.0.0.6 MAC address of VM E2:3A:DD:0A:8C:06 VM State RUNNING start-time 2015-05-07T18:22:44.136498+00:00 Last Reload Reason FIRST_BOOT Reboot Count 1</pre> <p>You must note the data plane location. In this example status of the data plane is shown under location 0/RP0/VM2.</p>
Start the data plane	<pre>sdr default-sdr location <data-plane-location> start</pre> <p>Example</p> <pre>sysadmin-vm:0_RP0# sdr default-sdr location 0/RP0/VM2 start Mon May 4 17:16:37.867 UTC Start ? [no,yes] yes result start sdr default-sdr location 0/RP0 request acknowledged</pre>
Shut the data plane down	<pre>sdr default-sdr location <data-plane-location> shut</pre> <p>Example</p> <pre>sysadmin-vm:0_RP0# sdr default-sdr location 0/RP0/VM2 shut Mon May 4 17:12:32.397 UTC Shut ? [no,yes] yes result shutdown sdr default-sdr location 0/RP0 request acknowledged</pre>

Task	Use this command
Reload the data plane	<p>sdr default-sdr location <data-plane-location> reload</p> <p>Example</p> <pre> sysadmin-vm:0_RP0# sdr default-sdr location 0/RP0/VM2 reload Mon May 4 17:21:17.390 UTC Reload ? [no,yes] yes result graceful reload sdr default-sdr location 0/RP0 request acknowledged </pre>



Note In the normal course of operations users must not manually start and stop the data plane manually.

Data Plane Debugging

Cisco IOS XRv 9000 Router provides a set of commands to check the status and statistics of the data plane, they are:



Note These data plane debug commands may briefly interrupt the traffic forwarding:

- show controller dpa statistics
- show controller dpa fib ipv4|ipv6 [<prefix> | summary]

- **show controller dpa version**—displays the version of the data plane.

For example:

```

RP/0/RP0/CPU0:ios#show controller dpa version
Fri May 29 19:28:16.520 UTC
Image built on 13:29:13 May 29 2015 in workspace /workspace1/shope/ttl_commit
DPA started May 29 18:11:23, up 0 days, 01:16

```

- **show controller dpa logging**—displays the data plane log. By default only errors and important events are log information is available.

For example:

```

RP/0/RP0/CPU0:SS_Node1#show controllers dpa logging
Mon Jun 29 19:47:33.245 UTC
Jun 29 01:43:32.820: Log File Started
Jun 29 01:43:32.820: DPA_INFO: DPA beginning initialization
Jun 29 01:43:32.823: DPA_INFO: Dataplane Agent enabled
Jun 29 01:43:32.823: DPA_INFO: Image built on 15:02:53 Jun 25 2015
Jun 29 01:43:32.823: DPA_INFO: Table WRED_STR of size 8388480 is being initialized
Jun 29 01:43:32.824: DPA_INFO: Table STATIC_POLICER_STR of size 8192 is being initialized
with data
Jun 29 01:43:32.824: DPA_INFO: Table HASH_DYN_BUCKET_STR of size 3355264 is being
initialized

```

```

Jun 29 01:43:32.825: DPA_INFO: Table HASH_BUCKET_STR of size 33552832 is being
initialized
Jun 29 01:43:32.829: DPA_INFO: Table DYN_FREE_BLOCK_STR of size 16777216 is being
initialized
Jun 29 01:43:32.832: DPA_INFO: Table INDEX_Q_STR of size 8192 is being initialized
Jun 29 01:43:32.832: DPA_INFO: Table HASH_HOST_DYN_BUCKET_STR of size 1677504 is being
initialized

```

- **show controller dpa statistics global**—displays the data plane statistics that includes drop packet counters, packet injected from the control plane and packet punted to the control plane.

For example:

```

RP/0/RP0/CPU0:ios#show control dpa statistics global
Fri May 29 19:27:35.497 UTC
Index  Punt                                     Count
-----
 1575  ARP                                           28
 1677  IFIB                                         1341
 1701  IPv4 FIB                                       5

Index  Inject                                     Count
-----
 267  IPv4 from fabric                               31
 268  IPv4 from fabric multicast                   1290
 270  IPv4 from fabric next-hop                     55
 275  Inject to fabric                             1376
 276  Inject to port                                11

Index  Drop                                     Count
-----
 63   Egress uIDB in down state                      1
 113  IPv6 disabled in uIDB                          52

```

- **show controller dpa fib ipv4|ipv6 [<prefix> | summary]**—displays FIB entries on the data plane. The <prefix> and summary keyword are optional.

For example:

```

RP/0/RP0/CPU0:ios#show controller dpa fib ipv4
Fri May 29 19:54:40.110 UTC

VRF id: 0
Default prefix 0.0.0.0/0 -> leaf:46423
total number of prefix:35
total_node_allocated:29 leaf_inserts:50 leaf_deletes 15 leaf_replaces: 2

Prefix          leaf_index
224.0.0.0/4     46436(0xb564)
224.0.0.0/24    46434(0xb562)
255.255.255.255/32 46429(0xb55d)
0.0.0.0/32     46433(0xb561)
10.1.1.1/32    46513(0xb5b1)
2.2.2.2/32     46510(0xb5ae)

```

If the summary keyword is used the command shows the total number of prefix in each vrf table and operational statistics. For example:

```

RP/0/RP0/CPU0:R1-PE1#show control dpa fib ipv4 summary

VRF id: 0
Default prefix 0.0.0.0/0 -> leaf:46421

```

```

total number of prefix: 27859
allocated nodes:      1089
leaf_inserts:        27922
leaf_deletes:        63
leaf_replaces:       173

VRF id: 1
Default prefix 0.0.0.0/0 -> leaf:46444
total number of prefix: 430
allocated nodes:      26
leaf_inserts:        430
leaf_deletes:        0
leaf_replaces:       213

```

If the <prefix> keyword is used the command displays list of all prefix and vrf tables that match the prefix. For example:

```

RP/0/RP0/CPU0:ios#show controller dpa fib ipv4 5.11.23.131/32
VRF id: 0
Prefix          leaf_index
5.11.23.131/32  1164818(0x11c612)

VRF id: 1

VRF id: 2

VRF id: 3

```

- **show controller dpa tm queue <num>**—displays the internal data of a traffic manager queue, including DRR weight, Q-limit, instantaneous packet and byte counts.

For example:

```

RP/0/RP0/CPU0:ios#show controller dpa tm queue 1
Fri May 29 19:31:25.556 UTC
Queue 1
  Parent Subport:    0
  Weight:            10
  Q-Limit:           625000
  Packets:           0
  Bytes:             0

```

- **show controller dpa tm subport <num>**—displays the internal data of a traffic manager subport, including DRR weight, shaped rate, queue config, and instantaneous packet.

For example:

```

RP/0/RP0/CPU0:ios#show controller dpa tm subport 3
Fri May 29 19:44:12.993 UTC
Subport 3
  Parent vPort:      3
  Weight:            10200
  Rate:              776726
  Being Deleted:     no
  Configured:        yes
  Queue 24 pkts:    0 bytes:    0
  Queue 25 pkts:    6 bytes:    8376
  Queue 26 pkts:    0 bytes:    0
  Queue 27 pkts:    0 bytes:    0
  Queue 28 pkts:    0 bytes:    0
  Queue 29 pkts:    0 bytes:    0

```

```
Queue    30 pkts:      0 bytes:      0
Queue    31 pkts:      0 bytes:      0
Priority Queues:      1
Best effort Queues:  7
```

- **show controller dpa tm vport <num>**—displays the internal data a traffic manager vport, including rate and whether or not flow control is currently active on that port.

For example:

```
RP/0/RP0/CPU0:ios#show controller dpa tm vport 0
Fri May 29 19:32:39.447 UTC
vPort 0
  Parent port:      0
  Rate:             95
  Flow control:     0
```