



Managing the Unicast RIB and FIB

This chapter describes how to manage routes in the unicast Routing Information Base (RIB) and the Forwarding Information Base (FIB) on the Cisco NX-OS switch.

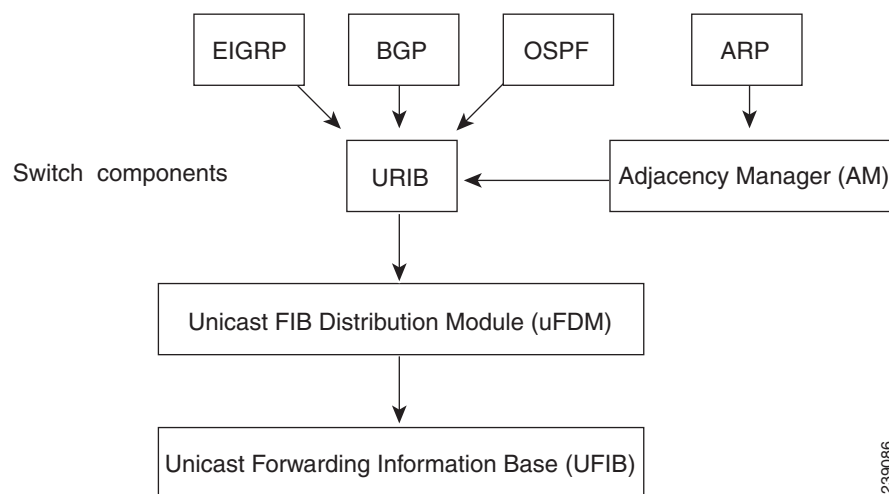
This chapter includes the following sections:

- [Information About the Unicast RIB and FIB, page 12-1](#)
- [Licensing Requirements for the Unicast RIB and FIB, page 12-2](#)
- [Managing the Unicast RIB and FIB, page 12-2](#)
- [Verifying the Unicast RIB and FIB Configuration, page 12-7](#)
- [Additional References, page 12-8](#)

Information About the Unicast RIB and FIB

The unicast RIB (IPv4 RIB) and FIB are part of the Cisco NX-OS forwarding architecture, as shown in [Figure 12-1](#).

Figure 12-1 Cisco NX-OS Forwarding Architecture



The unicast RIB maintains the routing table with directly connected routes, static routes, and routes learned from dynamic unicast routing protocols. The unicast RIB also collects adjacency information from sources such as the Address Resolution Protocol (ARP). The unicast RIB determines the best next hop for a given route and populates the unicast forwarding information base (FIBs) by using the services of the unicast FIB distribution module (FDM).

Each dynamic routing protocol must update the unicast RIB for any route that has timed out. The unicast RIB then deletes that route and recalculates the best next hop for that route (if an alternate path is available).

This section includes the following topic:

- [FIB Tables, page 12-2](#)

FIB Tables

The hardware provides two tables: a TCAM table and a hash table. The TCAM table is shared between longest prefix match (LPM) route /32 unicast route. The hash table is shared between the /32 unicast entries and the multicast entries. Each table has approximately 8000 routes.

If the LPM becomes 90% full, a warning messages appears. A message appears when there is sufficient space in the LPM and total usage is 70% or less. When the table is 100% full, the following message is displayed:

FIB_TCAM_RESOURCE_EXHAUSTION:FIB TCAM exhausted



Note

In order for the message to be displayed, the logging level for ipfib must be set to either 4 or 5.

Licensing Requirements for the Unicast RIB and FIB

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	The unicast RIB and FIB require no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .

Managing the Unicast RIB and FIB

This section includes the following topics:

- [Displaying Module FIB Information, page 12-3](#)
- [Configuring Load Sharing in the Unicast FIB, page 12-4](#)[Displaying Routing and Adjacency Information, page 12-4](#)
- [Clearing Forwarding Information in the FIB, page 12-5](#)[Estimating Memory Requirements for Routes, page 12-6](#)
- [Clearing Routes in the Unicast RIB, page 12-6](#)

**Note**

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Displaying Module FIB Information

You can display the FIB information on a switch.

DETAILED STEPS

To display the FIB information on a switch, use the following commands in any mode:

Command	Purpose
show ip fib adjacency { Example: switch# show ip fib adjacency	Displays the adjacency information for FIB.
show forwarding {ipv4 adjacency { Example: switch# show forwarding ipv4 adjacency	Displays the adjacency information for IPv4.
show ip fib interfaces Example: switch# show ip fib interfaces	Displays the FIB interface information for IPv4.
show ip fib route adjacency { Example: switch# show ip fib route	Displays the route table for IPv4.
show forwarding ipv4 route adjacency { Example: switch# show forwarding ipv4 route	Displays the route table for IPv4.

This example shows how to display the FIB contents on a switch:

```
switch# show ip fib route
```

```
IPv4 routes for table default/base
```

```
-----+-----+-----
Prefix      | Next-hop      | Interface
-----+-----+-----
0.0.0.0/32   | Drop          | Null0
255.255.255.255/32 | Receive      | sup-eth1
```

Configuring Load Sharing in the Unicast FIB

Dynamic routing protocols, such as Open Shortest Path First (OSPF), support load balancing with equal-cost multipath (ECMP). The routing protocol determines its best routes based on the metrics configured for the protocol and installs up to the protocol-configured maximum paths in the unicast RIB. The unicast RIB compares the administrative distances of all routing protocol paths in the RIB and selects a best path set from all of the path sets installed by the routing protocols. The unicast RIB installs this best path set into the FIB for use by the forwarding plane.

The forwarding plane uses a load-sharing algorithm to select one of the installed paths in the FIB to use for a given data packet.

You can globally configure the following load-sharing settings:

- **load-share mode**—Selects the best path based on the destination address and port or the source and the destination address and port.
- **Universal ID**—Sets the random seed for the hash algorithm. You do not need to configure the Universal ID. Cisco NX-OS chooses the Universal ID if you do not configure it.

**Note**

Load sharing uses the same path for all packets in a given flow. A flow is defined by the load-sharing method that you configure. For example, if you configure source-destination load sharing, then all packets with the same source IP address and destination IP address pair follow the same path.

To configure the unicast FIB load-sharing algorithm, use the following command in global configuration mode:

Command	Purpose
ip load-sharing address { destination port destination source-destination [port source-destination] } [universal-id seed]	Configures the unicast FIB load-sharing algorithm for data traffic. The <i>universal-id</i> range is from 1 to 4294967295.
Example: switch(config)# ip load-sharing address source-destination	

To display the unicast FIB load-sharing algorithm, use the following command in any mode:

Command	Purpose
show ip load-sharing	Displays the unicast FIB load-sharing algorithm for data traffic.
Example: switch(config)# show ip load-sharing	

Displaying Routing and Adjacency Information

You can display the routing and adjacency information.

To display the routing and adjacency information, use the following commands in any mode:

Command	Purpose
show ip route [<i>route-type</i> interface <i>int-type number</i> next-hop] Example: switch# show ip route	Displays the unicast route table. The <i>route-type</i> argument can be a single route prefix, direct, static, or a dynamic route protocol. Use the ? keyword to see the supported interfaces.
show ip adjacency [<i>prefix</i> <i>interface-type number</i> [summary] non-best] [detail] [vrf <i>vrf-id</i>] Example: switch# show ip adjacency	Displays the adjacency table. The argument ranges are as follows: <ul style="list-style-type: none"> <i>prefix</i>—Any IPv4 prefix address. <i>interface-type number</i>—Use the ? keyword to see the supported interfaces. <i>vrf-id</i>—Any case-sensitive, alphanumeric string up to 32 characters.
show ip routing [<i>route-type</i> interface <i>int-type number</i> next-hop recursive-next-hop summary updated { since until } <i>time</i>] Example: switch# show routing summary	Displays the unicast route table. The <i>route-type</i> argument can be a single route prefix, direct, static, or a dynamic route protocol. Use the ? keyword to see the supported interfaces.

This example shows how to display the unicast route table:

```
switch# show ip route
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
*** denotes best mcast next-hop
'[x/y]' denotes [preference/metric]

192.168.0.2/24, ubest/mbest: 1/0, attached
    *via 192.168.0.32, Eth1/5, [0/0], 22:34:09, direct
192.168.0.32/32, ubest/mbest: 1/0, attached
    *via 192.168.0.32, Eth1/5, [0/0], 22:34:09, local
```

This example shows the adjacency information:

```
switch# show ip adjacency

IP Adjacency Table for VRF default
Total number of entries: 2
Address           Age           MAC Address    Pref Source    Interface    Best
10.1.1.1          02:20:54     00e0.b06a.71eb  50  arp        mgmt0        Yes
10.1.1.253        00:06:27     0014.5e0b.81d1  50  arp        mgmt0        Yes
```

Clearing Forwarding Information in the FIB

You can clear one or more entries in the FIB. Clearing a FIB entry does not affect the unicast RIB.



Caution

The **clear forwarding** command disrupts forwarding on the switch.

To clear an entry in the FIB, including a Layer 3 inconsistency, use the following command in any mode:

Estimating Memory Requirements for Routes

Command	Purpose
clear forwarding { ipv4 } route { * <i>prefix</i> } [vrf <i>vrf-name</i>] [module { <i>slot</i> all }] Example: switch(config)# clear forwarding ipv4 route *	Clears one or more entries from the FIB. The route options are as follows: <ul style="list-style-type: none"> *—All routes. <i>prefix</i>—Any IP prefix. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters. The <i>slot</i> range is from 1 to 10.

You can estimate the memory that a number of routes and next-hop addresses will use.

To estimate the memory requirements for routes, use the following command in any mode:

Command	Purpose
show routing memory estimate routes <i>num-routes</i> next-hops <i>num-nexthops</i> Example: switch# show routing memory estimate routes 1000 next-hops 1	Displays the memory requirements for routes. The <i>num-routes</i> range is from 1000 to 1000000. The <i>num-nexthops</i> range is from 1 to 16.

Clearing Routes in the Unicast RIB

You can clear one or more routes from the unicast RIB.



Caution

The * keyword is severely disruptive to routing.

To clear one or more entries in the unicast RIB, use the following commands in any mode:

Command	Purpose
clear ip route <i>{* {route prefix/length} [next-hop interface]}</i> <i>[vrf vrf-name]</i> Example: switch(config)# clear ip route 10.2.2.2	<p>Clears one or more routes from both the unicast RIB and all the module FIBs. The route options are as follows:</p> <ul style="list-style-type: none"> • <i>*</i>—All routes. • <i>route</i>—An individual IP route. • <i>prefix/length</i>—Any IP prefix. • <i>next-hop</i>—The next-hop address • <i>interface</i>—The interface to reach the next-hop address. <p>The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.</p>
clear routing <i>[multicast unicast]</i> <i>[ip ipv4] {* {route prefix/length} [next-hop interface]}</i> <i>[vrf vrf-name]</i> Example: switch(config)# clear routing ip 10.2.2.2	<p>Clears one or more routes from the unicast RIB. The route options are as follows:</p> <ul style="list-style-type: none"> • <i>*</i>—All routes. • <i>route</i>—An individual IP route. • <i>prefix/length</i>—Any IP prefix. • <i>next-hop</i>—The next-hop address • <i>interface</i>—The interface to reach the next-hop address. <p>The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 32 characters.</p>

Verifying the Unicast RIB and FIB Configuration

To display the unicast RIB and FIB configuration information, perform one of the following tasks:

Command	Purpose
show forwarding adjacency	Displays the adjacency table on a module.
show forwarding distribution <i>{ clients fib-state }</i>	Displays the FIB distribution information.
show forwarding interfaces <i>module slot</i>	Displays the FIB information for a module.
show forwarding ipv4 route	Displays routes in the FIB.
show ip adjacency	Displays the adjacency table.
show ip route	Displays IPv4 routes from the unicast RIB.
show routing	Displays routes from the unicast RIB.

Additional References

For additional information related to managing unicast RIB and FIB, see the following sections:

- [Related Documents, page 12-8](#)

Related Documents

Related Topic	Document Title
Unicast RIB and FIB CLI commands	<i>Cisco Nexus 5000 Series Command Reference, Cisco NX-OS Releases 4.x, 5.x</i>