



Configuring the Unicast RIB and FIB

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

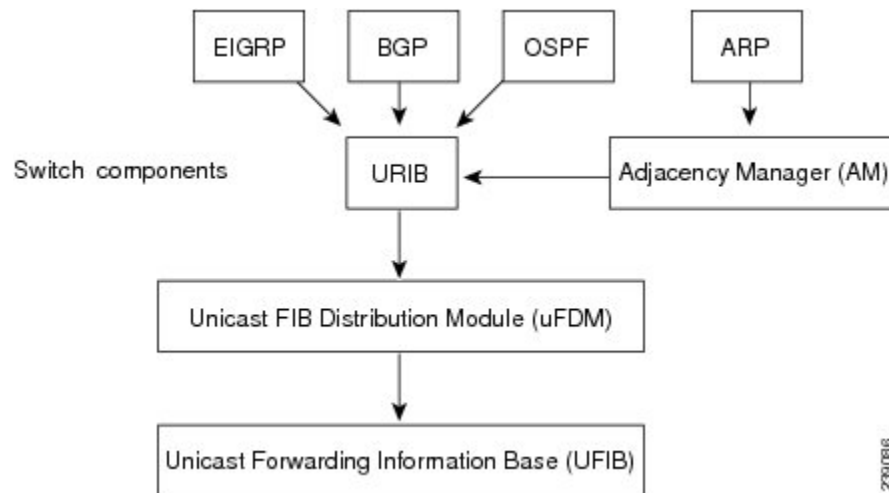
This chapter includes the following sections:

- [About the Unicast RIB and FIB, on page 1](#)
- [Managing the Unicast RIB and FIB, on page 3](#)
- [Verifying the Unicast RIB and FIB Configuration, on page 11](#)

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 the following figure.

Figure 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).

Layer 3 Consistency Checker

In rare instances, an inconsistency can occur between the unicast RIB and the FIB on each module. Cisco NX-OS supports the Layer 3 consistency checker. This feature detects inconsistencies between the unicast IPv4 RIB and the FIB on each interface module. Inconsistencies include the following:

- Missing prefix
- Extra prefix
- Wrong next-hop address
- Incorrect Layer 2 rewrite string in the ARP or neighbor discovery (ND) cache

The Layer 3 consistency checker compares the FIB entries to the latest adjacency information from the Adjacency Manager (AM) and logs any inconsistencies. The consistency checker then compares the unicast RIB prefixes to the module FIB and logs any inconsistencies. See the [Triggering the Layer 3 Consistency Checker](#) section.

You can then manually clear any inconsistencies. See the [Clearing Forwarding Information in the FIB](#) section.

FIB Tables

The following are the unicast routing table capacities for the Cisco Nexus 3500 platform switches when the switch is configured in the normal forwarding mode:

- Unicast Routing Host table = 64,000 hash table entries
- Unicast Routing LPM table = 16,000 TCAM entries
- ECMP members Table size: 16,000 entries

The following are the unicast routing table capacities for the Cisco Nexus 3500 platform switches when the switch is configured in the warp mode:

- L3 Unicast Host table = 8000 TCAM entries
- L3 Unicast LPM table = 4000 TCAM entries



Note ECMP is not supported on warp mode.



Note In warp mode, when two equal cost paths are received in RIB, one of the paths is installed in the hardware. It is recommended to configure the maximum-path to one under the routing protocol configuration.

Virtualization Support

The Unicast RIB and FIB support Virtual Routing and Forwarding instances (VRFs). For more information, see [Configuring Layer 3 Virtualization](#).

Managing the Unicast RIB and FIB



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: <pre>switch# show ip fib adjacency</pre> | Displays the adjacency information for FIB. |
| show forwarding ipv4 adjacency Example: <pre>switch# show forwarding ipv4 adjacency</pre> | Displays the adjacency information for IPv4. |
| show ip fib interfaces Example: <pre>switch# show ip fib interfaces</pre> | Displays the FIB interface information for IPv4. |
| show ip fib route Example: <pre>switch# show ip fib route</pre> | Displays the route table for IPv4. |
| show forwarding ipv4 route Example: <pre>switch# show forwarding ipv4 route</pre> | Displays the route table for IPv4. |

This example shows 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.

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 |
|--|---|
| <p>ip load-sharing address {destination port destination source-destination [port source-destination] source } hardware lb-keyshift value lb-2nd-heir-keyshift value [universal-id seed] [rotate rotate] [concatenation]</p> <p>Example:</p> <pre>ip load-sharing address source-destination port source-destination hardware lb-keyshift 1 lb-2nd-hier-keyshift 10</pre> | <p>Configures the unicast FIB load-sharing algorithm for data traffic.</p> <p>Beginning with Cisco NX-OS Release 10.3(3)F, the hardware keyword is added to support the following parameters in the IHB_ECMP_LB_KEY_CFG tables only on Cisco Nexus 9600-R/RX line cards:</p> <ul style="list-style-type: none"> • lb-keyshift: Sets the ECMP_LB_KEY_SHIFT value for load balancing. The range is 1-10. • lb-2nd-hier-keyshift: Sets the ECMP_2ND_HIER_LB_KEY_SHIFT value for load balancing. The range is 1-10. <p>The following options are available for all IP load sharing configurations:</p> <ul style="list-style-type: none"> • The universal-id option sets the random seed for the hash algorithm and shifts the flow from one link to another. <p>You do not need to configure the universal ID. Cisco NX-OS chooses the universal ID if you do not configure it. The <i>universal-id</i> range is from 1 to 4294967295.</p> <ul style="list-style-type: none"> • The rotate option causes the hash algorithm to rotate the link picking selection so that it does not continually choose the same link across all nodes in the network. It does so by influencing the bit pattern for the hash algorithm. This option shifts the flow from one link to another and load balances the already load-balanced (polarized) traffic from the first ECMP level across multiple links. <p>If you specify a <i>rotate</i> value, the 64-bit stream is interpreted starting from that bit position in a cyclic rotation. The <i>rotate</i> range is from 1 to 63, and the default is 32.</p> <p>Note With multi-tier Layer 3 topology, polarization is possible. To avoid polarization, use a different rotate bit at each tier of the topology.</p> <p>Note To configure a rotation value for port channels, use the port-channel load-balance src-dst ip-l4port rotate rotate command. For more information on this command, see the <i>Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide</i>.</p> <ul style="list-style-type: none"> • The concatenation option ties together the hash tag values for ECMP and the hash tag values for port channels in order to use a stronger 64-bit hash. If you do not use this option, you can control ECMP load-balancing and port-channel load-balancing independently. The default is disabled. |

Configuring Hash Offset

To avoid ECMP polarization in a multi-tier ECMP session, you must configure a different ECMP hash-offset on each tier. Starting with Cisco NX-OS Release 6.0(2)U5(1), a new CLI for ECMP hash

concatenation is introduced to achieve uniform distribution of the traffic across 16 way ECMP paths. The updated CLI support exists for Cisco Nexus 3100 platform switches and not on Cisco Nexus 3000 Series switches. You can configure the hash offset in the range of <0-15> in non-concatenate mode and in the range of <0-63> in concatenate mode.

In concatenate mode, if the hash-offset is set to 0 and concatenation is set, the **show running-config** command displays hardware ecmp hash-offset 0 concatenate. The hash-offset is programmed as per value. On downgrade, if concatenation is configured, the CAP check asks to remove the configuration.

In non-concatenate mode, if the hash-offset is set to 0 and concatenation is reset, the show running-config command does not display hardware ecmp hash-offset 0. The hash-offset is programmed as per value if the hash-offset value is in range 0-15. The hash-offset displays CLI error if the value is in range 16-63 (The non-concatenated mode supports hash-offset for 0-15 range).

The hash-offset in show running-config is visible as per the configured value.

To configure an ECMP hash-offset, use the following commands in global configuration mode:

| Command | Purpose |
|---|--|
| # hardware ecmp hash-offset ? | Configures the ECMP hash-offset. The range is from 0 to 63. The hash offset in the range <0-15> is for the non-concatenate mode. The hash offset in the range <0-63> is for the concatenate mode. |
| #hardware ecmp hash-offset <i>number</i> Example: switch(config)# hardware ecmp hash-offset 5 | Configures the ECMP hash-offset in the non-concatenate mode. The range is from 0 to 15. The default value is 0. |
| #hardware ecmp hash-offset <0-63> concatenate Example: switch(config)# hardware ecmp hash-offset 63 concatenate | Configures the ECMP hash-offset in the concatenate mode. Note Concatenation support exists for Cisco Nexus 3100 platform switches and not for Cisco Nexus 3000 Series switches. This CLI generates an error on Cisco Nexus 3000 Series switches. |

You can use the ECMP hash-offset configured by using the hardware ecmp hash-offset command along with different universal IDs configured by using the ip load-sharing address command to produce various hash results in the load-sharing algorithm.

Configuring Hash Polynomial

Starting with Release 6.0(2)U5(1), new CLI is added for the CRC configuration.

| Command | Purpose |
|---|--|
| switch# config t | Enters configuration mode. |
| switch(config)# hardware ecmp ? hash-offset Configure hash offsethash-polynomial Configure hash polynomial | Displays hash-offset and hash-polynomial as the configuration options for hardware ECMP. |

| Command | Purpose |
|---|--|
| switch config)# hardware ecmp hash-polynomial ? CRC16 Hash polynomial CRC16CRC32HI Hash polynomial CRC32 HI | Displays CRC16 and CRC32HI as configuration options for Hash polynomial. |
| switch config)#show running-config | Displays the running configuration. |

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 | |

This example shows the output of **show ip load-sharing** command:

```
hardware lb-keyshift 1 lb-2nd-hier-keyshift 10
switch(config)# ip load-sharing address source-destination port source-destination
switch(config)# show ip load-sharing
IPv4/IPv6 ECMP load sharing:
Universal-id (Random Seed): 251533739
Load-share mode : address source-destination port source-destination
GRE-Outer hash is disabled
Concatenation is disabled
Rotate: 32

Lbkeyshift: 1
2ndHeirLbkeyshift: 10
switch(config)#
```

To display the route that the unicast RIB and FIB use for a particular source address and destination address, use the following command in any mode:

| Command | Purpose |
|---|--|
| show routing hash <i>source-addr dest-addr</i> [ip-proto <i>ip-protocol</i>] [<i>source-l4-port</i> <i>dest-l4-port</i>] [vrf <i>vrf-name</i>] | Displays the route that the unicast RIB FIB use for a source and destination address pair. The source address and destination address format is x.x.x.x. The source port and destination port range is from 1 to 65535. The VRF name can be any case-sensitive, alphanumeric string up to 64 characters. The ip-proto option corresponds to the protocol field of the IP header. |
| Example: switch# show routing hash 1.1.1.6.5.5 5.3 ip-proto 0x11 10 234 | |

This example shows how to display the route selected for a source/destination pair:

```
switch# show routing hash 1.1.1.6.5.5.5.3 ip-proto 0x11 10 234
Load-share parameters used for software forwarding:
load-share mode: address source-destination port source-destination
Universal-id seed: 0xe05e2e85
Invoking pc_ic_ecmp_resolution
Hash for VRF "default"
Hashing to path *Eth1/29%
For route:
5.5.5.0/24 ubest/mbest: 3/0
*via 2.2.2.1, Eth1/18, [1/0], 00:14:14, static
```

```
*via 3.3.3.1, Eth1/29, [1/0], 00:14:14, static
*via 4.4.4.1, Eth1/34, [1/0], 00:14:14, static
```

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</i> number next-hop] Example: <pre>switch# show ip route</pre> | 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 number</i> [summary] non-best] [detail] [vrf <i>vrf-id</i>] Example: <pre>switch# show ip adjacency</pre> | Displays the adjacency table. The argument ranges are as follows: <ul style="list-style-type: none"> • <i>prefix</i> —Any IPv4prefix 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</i> number next-hop recursive-next-hop summary updated { since until } <i>time</i>] Example: <pre>switch# show routing summary</pre> | 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 displays 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
```


Triggering the Layer 3 Consistency Checker

You can manually trigger the Layer 3 consistency checker.

To manually trigger the Layer 3 consistency checker, use the following commands in global configuration mode:

| Command | Purpose |
|---|--|
| test [ipv4] [unicast] forwarding inconsistency [vrf vrf-name] [module { slot all }] Example: <pre>switch(config)# test forwarding inconsistency</pre> | Starts a Layer 3 consistency check. 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. |

To stop the Layer 3 consistency checker, use the following commands in global configuration mode:

| Command | Purpose |
|---|---|
| test forwarding [ipv4] [unicast] inconsistency [vrf vrf-name] [module { slot all }] stop Example: <pre>switch(config)# test forwarding inconsistency stop</pre> | Stops a Layer 3 consistency check. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>slot</i> range is from 1 to 10. |

To display the Layer 3 inconsistencies, use the following commands in any mode:

| Command | Purpose |
|---|---|
| show forwarding [ipv4] inconsistency [vrf vrf-name] [module { slot all }] Example: <pre>switch(config)# show forwarding inconsistency</pre> | Displays the results of a Layer 3 consistency check. 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. |

Clearing Forwarding Information in the FIB

You can clear one or more entries in the FIB.



Note 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:

| Command | Purpose |
|---|---|
| <p>clear forwarding { ipv4 } route { * <i>prefix</i> } [vrf <i>vrf-name</i>] [module { <i>slot</i> all }]</p> <p>Example:</p> <pre>switch(config)# clear forwarding ipv4 route *</pre> | <p>Clears one or more entries from the FIB. The route options are as follows:</p> <ul style="list-style-type: none"> • *—All routes. • <i>prefix</i> —Any IPprefix. <p>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.</p> |



Note Ensure you clear the RIB entry after you clear the FIB entry.

Estimating Memory Requirements for Routes

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 |
|---|---|
| <p>show routing memory estimate routes <i>num-routes</i> next-hops <i>num-nexthops</i></p> <p>Example:</p> <pre>switch# show routing memory estimate routes 1000 next-hops 1</pre> | <p>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.</p> |

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 |
|---|---|
| <p>clear iproute <i>{* { route prefix/length } [next-hop interface]} [vrf vrf-name]</i></p> <p>Example:</p> <pre>switch(config)# clear ip route 10.2.2.2</pre> | <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 IPRoute. • <i>prefix/length</i> —Any IPprefix. • <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> |
| <p>clear routing unicast [<i>ip ipv4</i>] <i>{* { route prefix/length } [next-hop interface]} [vrf vrf-name]</i></p> <p>Example:</p> <pre>switch(config)# clear routing ip 10.2.2.2</pre> | <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 IPRoute. • <i>prefix/length</i> —Any IPprefix. • <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 module <i>slot</i> | Displays the FIB information for a module. |
| show forwarding ipv4route | Displays routes in the FIB. |
| show ip adjacency | Displays the adjacency table. |
| show ip route | Displays IPv4routes from the unicast RIB. |
| show routing | Displays routes from the unicast RIB. |

