



CHAPTER 14

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 device.

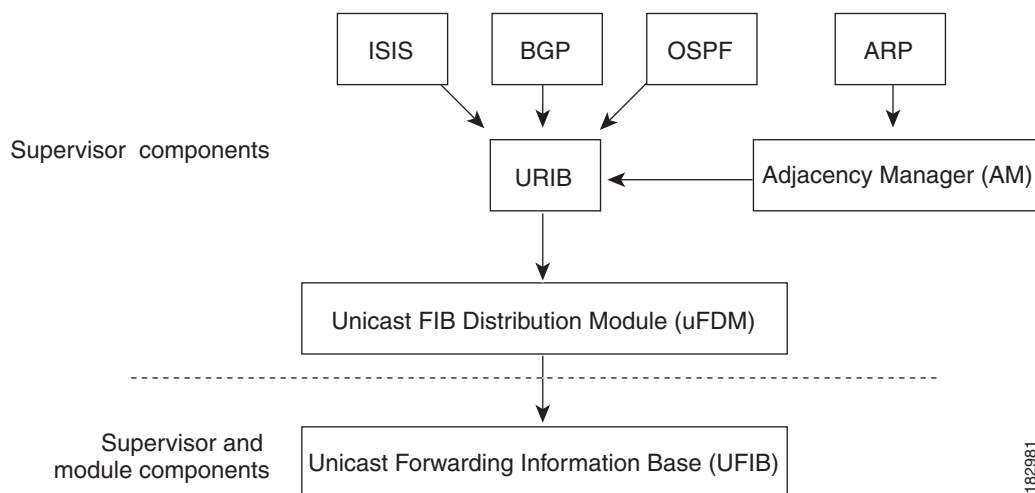
This chapter includes the following sections:

- [Information About the Unicast RIB and FIB, page 14-1](#)
- [Licensing Requirements for the Unicast RIB and FIB, page 14-2](#)
- [Guidelines and Limitations, page 14-3](#)
- [Managing the Unicast RIB and FIB, page 14-3](#)
- [Verifying the Unicast RIB and FIB, page 14-9](#)
- [Additional References, page 14-10](#)
- [Feature History for Unicast RIB and FIB, page 14-10](#)

Information About the Unicast RIB and FIB

The unicast RIB (IPv4 RIB and IPv6 RIB) and FIB are part of the Cisco NX-OS forwarding architecture, as shown in [Figure 14-1](#).

Figure 14-1 Cisco NX-OS Forwarding Architecture



[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

The unicast RIB exists on the active supervisor. It 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 bases (FIB) on the modules 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 topics:

- [Layer 3 Consistency Checker, page 14-2](#)
- [Virtualization Support, page 14-2](#)

Layer 3 Consistency Checker

In rare instances, an inconsistency can occur between the unicast RIB and the FIB on each module. In Cisco NX-OS Release 4.0(3) and later releases, Cisco NX-OS supports the Layer 3 consistency checker. This feature detects inconsistencies between the unicast IPv4 RIB on the supervisor module 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 on page 14-7](#).

You can then manually clear any inconsistencies. See the [“Clearing Forwarding Information in the FIB” section on page 14-8](#).

Virtualization Support

The Unicast RIB and FIB support Virtual Routing and Forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco NX-OS Virtual Device Context Configuration Guide*, and see [Chapter 13, “Configuring Layer 3 Virtualization.”](#)

Licensing Requirements for the Unicast RIB and FIB

The following table shows the licensing requirements for this feature:

Send document comments to nexus7k-docfeedback@cisco.com.

Product	License Requirement
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 NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide, Release 4.0</i> .

Guidelines and Limitations

Unicast RIB and FIB have the following guidelines and limitations and restrictions:

- The FIB TCAM is 128,000 physical entries, which is divided into the following sections by default:
 - 56,000 IPv4 unicast routes (56,000 physical entries)
 - 32,000 IPv4 multicast routes or IPv6 unicast routes (64,000 physical entries)
 - 2000 IPv6 multicast routes (8000 physical entries)
- You cannot change the default FIB division.

Managing the Unicast RIB and FIB

This section includes the following topics:

- [Displaying Module FIB Information, page 14-4](#)
- [Configuring Load Sharing in the Unicast RIB, page 14-4](#)
- [Displaying Routing and Adjacency Information, page 14-6](#)
- [Triggering the Layer 3 Consistency Checker, page 14-7](#)
- [Clearing Forwarding Information in the FIB, page 14-8](#)
- [Clearing Routes in the Unicast RIB, page 14-9](#)



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.

Send document comments to nexus7k-docfeedback@cisco.com.

Displaying Module FIB Information

You can display the FIB information on a module.

DETAILED STEPS

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

Command	Purpose
show ip fib adjacency module slot Example: switch# show ip fib adjacency module 2	Displays the adjacency information for IPv4.
show forwarding {ipv4 ipv6} adjacency module slot Example: switch# show forwarding ipv6 adjacency module 2	Displays the adjacency information for IPv4 or IPv6.
show ip fib interfaces module slot Example: switch# show ip fib interfaces module 2	Displays the FIB interface information for IPv4.
show ip fib route module slot Example: switch# show ip fib route module 2	Displays the route table for IPv4.
show forwarding {ipv4 ipv6} route module slot Example: switch# show forwarding ipv6 route module 2	Displays the route table for IPv4 or IPv6.

This example shows the FIB contents on a module:

```
switch# show ip fib route module 2
```

```
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 RIB

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.

Send document comments to nexus7k-docfeedback@cisco.com.

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 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.
- **Hash seed**—Sets the random seed for the hash algorithm.

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

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

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

Command	Purpose
<pre>show ip load-sharing</pre> <p>Example: switch(config)# show ip load-sharing address source-destination</p>	Configures the unicast RIB load-sharing algorithm for data traffic. The hash seed range is from 1 to 4294967295.

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

Command	Purpose
<pre>show routing hash source-addr dest-addr [source-port dest-port] [vrf vrf-name]</pre> <p>Example: switch# show routing hash 192.0.2.1 10.0.0.1</p>	Configures the unicast RIB load-sharing algorithm for data traffic. 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.

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

```
switch# show routing hash 10.0.0.5 30.0.0.2
Load-share parameters used for software forwarding:
load-share mode: address source-destination port source-destination
Universal-id seed: 0xe05e2e85
Hash for VRF "default"
Hashing to path *20.0.0.2 (hash: 0x0e), for route:
```

Send document comments to nexus7k-docfeedback@cisco.com.

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
<pre>show {ip ipv6} route [<i>route-type</i> interface <i>int-type number</i> next-hop]</pre> <p>Example: switch# show ip route</p>	<p>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 ? command to see the supported interfaces.</p>
<pre>show {ip ipv6} adjacency [<i>prefix</i> <i>interface-type number</i> [summary] non-best] [detail] [vrf <i>vrf-id</i>]</pre> <p>Example: switch# show ip adjacency</p>	<p>Displays the adjacency table. The argument ranges are as follows:</p> <ul style="list-style-type: none"> <i>prefix</i>—Any IPv4 or IPv6 prefix address. <i>interface-type number</i>—Use the ? command to see the supported interfaces. <i>vrf-id</i>—Any case-sensitive alphanumeric string up to 64 characters.
<pre>show {ip ipv6} routing [<i>route-type</i> interface <i>int-type number</i> next-hop recursive-next-hop summary updated {since until} <i>time</i>]</pre> <p>Example: switch# show routing summary</p>	<p>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 ? command to see the supported interfaces.</p>

This example displays the unicast route table:

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

0.0.0.0/0, 1 ucast next-hops, 0 mcast next-hops
  *via 10.1.1.1, mgmt0, [1/0], 5d21h, static
0.0.0.0/32, 1 ucast next-hops, 0 mcast next-hops
  *via Null0, [220/0], 1w6d, local, discard
10.1.0.0/22, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.1.55, mgmt0, [0/0], 5d21h, direct
10.1.0.0/32, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.0.0, Null0, [0/0], 5d21h, local
10.1.1.1/32, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.1.1, mgmt0, [2/0], 5d16h, am
10.1.1.55/32, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.1.55, mgmt0, [0/0], 5d21h, local
10.1.1.253/32, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.1.253, mgmt0, [2/0], 5d20h, am
10.1.3.255/32, 1 ucast next-hops, 0 mcast next-hops, attached
  *via 10.1.3.255, mgmt0, [0/0], 5d21h, local
255.255.255.255/32, 1 ucast next-hops, 0 mcast next-hops
  *via Eth Inband Port, [0/0], 1w6d, local
```

Send document comments to nexus7k-docfeedback@cisco.com.

This example shows the adjacency information:

```
switch# show ip adjacency

IP Adjacency Table for context 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
switch
```

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
<pre>test forwarding inconsistency [ip ip4] [unicast] [vrf vrf-name] [module {slot all}]</pre> <p>Example: switch(config)# test forwarding inconsistency</p>	<p>Starts 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.</p>

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

Command	Purpose
<pre>test forwarding inconsistency [ip ip4] [unicast] [vrf vrf-name] [module {slot all}] stop</pre> <p>Example: switch(config)# test forwarding inconsistency stop</p>	<p>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.</p>

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

Command	Purpose
<pre>show forwarding inconsistency [ip ip4] [unicast] [vrf vrf-name] [module {slot all}]</pre> <p>Example: switch(config)# show forwarding inconsistency</p>	<p>Displays the results of 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.</p>

Send document comments to nexus7k-docfeedback@cisco.com.

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 device.

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

Command	Purpose
<pre>clear forwarding {ip ipv4 ipv6} route {* prefix} [vrf vrf-name] [module {slot all}]</pre> <p>Example: switch(config)# clear forwarding ip route *</p>	<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 IP or IPv6 prefix. <p>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.</p>

Send document comments to nexus7k-docfeedback@cisco.com.

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
<pre>clear {ip ipv4 ipv6} route {* {route prefix/length}[next-hop interface]} [vrf vrf-name]</pre> <p>Example: switch(config)# clear ip route 10.2.2.2</p>	<p>Clears one or more routes from both the unicast RIB and the all the module FIBs. The route options are as follows:</p> <ul style="list-style-type: none"> <code>*</code>—All routes. <code>route</code>—An individual IP or IPv6 route. <code>prefix/length</code>—Any IP or IPv6 prefix. <code>next-hop</code>—The next-hop address <code>interface</code>—The interface to reach the next-hop address. <p>The <code>vrf-name</code> can be any case-sensitive alphanumeric string up to 64 characters.</p>
<pre>clear routing [multicast unicast] [ip ipv4 ipv6] {* {route prefix/length}[next-hop interface]} [vrf vrf-name]</pre> <p>Example: switch(config)# clear routing ip 10.2.2.2</p>	<p>Clears one or more routes from the unicast RIB. The route options are as follows:</p> <ul style="list-style-type: none"> <code>*</code>—All routes. <code>route</code>—An individual IP or IPv6 route. <code>prefix/length</code>—Any IP or IPv6 prefix. <code>next-hop</code>—The next-hop address <code>interface</code>—The interface to reach the next-hop address. <p>The <code>vrf-name</code> can be any case-sensitive alphanumeric string up to 64 characters.</p>

Verifying the Unicast RIB and FIB

To verify the unicast RIB and FIB information, use the following commands:

Command	Purpose
<code>show forwarding adjacency</code>	Displays the adjacency table on a module.
<code>show forwarding distribution {clients fib-state}</code>	Displays the FIB distribution information.
<code>show forwarding interfaces module slot</code>	Displays the FIB information for a module.
<code>show forwarding {ip ipv4 ipv6} route</code>	Displays routes in the FIB.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
<code>show {ip ipv6} adjacency</code>	Displays the adjacency table.
<code>show {ip ipv6} route</code>	Displays IPv4 or IPv6 routes from the unicast RIB.
<code>show routing</code>	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 14-10](#)
- [Feature History for Unicast RIB and FIB, page 14-10](#)

Related Documents

Related Topic	Document Title
unicast RIB and FIB CLI commands	<i>Cisco NX-OS Unicast Routing Command Reference, Release 4.0</i>

Feature History for Unicast RIB and FIB

[Table 14-1](#) lists the release history for this feature.

Table 14-1 *Feature History for Unicast RIB and FIB*

Feature Name	Releases	Feature Information
Unicast RIB and FIB	4.0(3)	Added support to clear individual routes in unicast RIB and FIB.
Unicast RIB and FIB	4.0(1)	This feature was introduced.