



RIB User Guide

Version 24.1.1, 31 January 2024

Registered Address	Support	Sales
26, Kingston Terrace, Princeton, New Jersey 08540, United States		
		+91 80 4850 5445
http://www.rtbrick.com	support@rtbrick.com	sales@rtbrick.com

©Copyright 2024 RtBrick, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos and service marks ("Marks") displayed in this documentation are the property of RtBrick in the United States and other countries. Use of the Marks are subject to RtBrick's Term of Use Policy, available at <https://www.rtbrick.com/privacy>. Use of marks belonging to other parties is for informational purposes only.

Table of Contents

1. Overview	3
1.1. Routing Instances	3
1.2. Route Preference	4
1.3. Route Selection	5
1.4. Route Resolution	5
1.5. Policy Attachments	6
1.6. Nexthops	6
1.7. Adjacency	6
1.8. FIB	7
2. Operational commands	8
2.1. Display Routes	8
2.2. Displaying Route Resolution	11

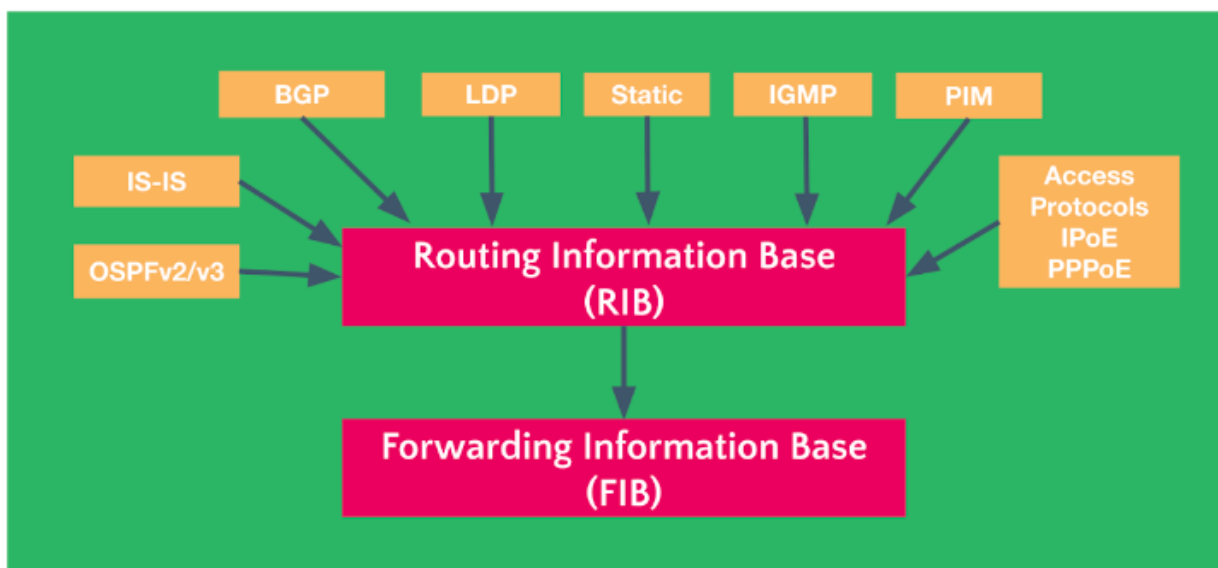
1. Overview

A Routing Information Base (RIB) stores and maintains the route information. The RIB table contains information about each reachable network prefix and the next hop information. Each routing protocol inserts its routes into the RIB when it learns a new route. The RBFS route manager (ribd) selects the best routes from the protocol RIBs and places them into the Forwarding Information Base (FIB) table. If a destination becomes unreachable, the route is marked unusable and eventually removed from the RIB.

The main differences between RIB and FIB are:


- RIB is about storing and maintaining route information, including attributes for route selection (control plane), etc.
- FIB is used to determine how packets are forwarded (data plane), i.e., only contains a subset of information from RIB. Also, FIB is downloaded to ASIC or VPP.

The image below shows how various routing protocols insert their routes into the RIB.



1.1. Routing Instances

A routing instance represents a set of interfaces, routing protocols, and corresponding routing tables. You can think of a routing instance as a virtual router within RBFS.



Each logical interface (ifl) is associated uniquely with a routing instance. If no routing instance is explicitly configured, the default routing instance is used. The interface association can be seen with the `show interface address` command.

Routing instances are configured using the `set instance <instance_name>` syntax. When configuring a routing instance, you must specify which address families (AFI, SAF) the routing instance should support.

1.2. Route Preference

Several routing protocols, including static routes, may provide multiple paths to reach a particular destination. However, not all of these paths are necessarily optimal. At any given time, only one routing protocol can identify the most optimal route to the destination. Each routing protocol, including static routes, is given a preference value, where lower values indicate higher preference. In cases where multiple routing sources are present, the route identified by the routing protocol with the highest preference is selected as the best option and added to the local routing table.

Routing Protocol	Route Preference
Direct	0
Static	2
IPoE	7
PPP	8
LDP	9
OSPF	10
IS-IS LEVEL 1	15
IS-IS LEVEL 2	18
eBGP	20
iBGP	200

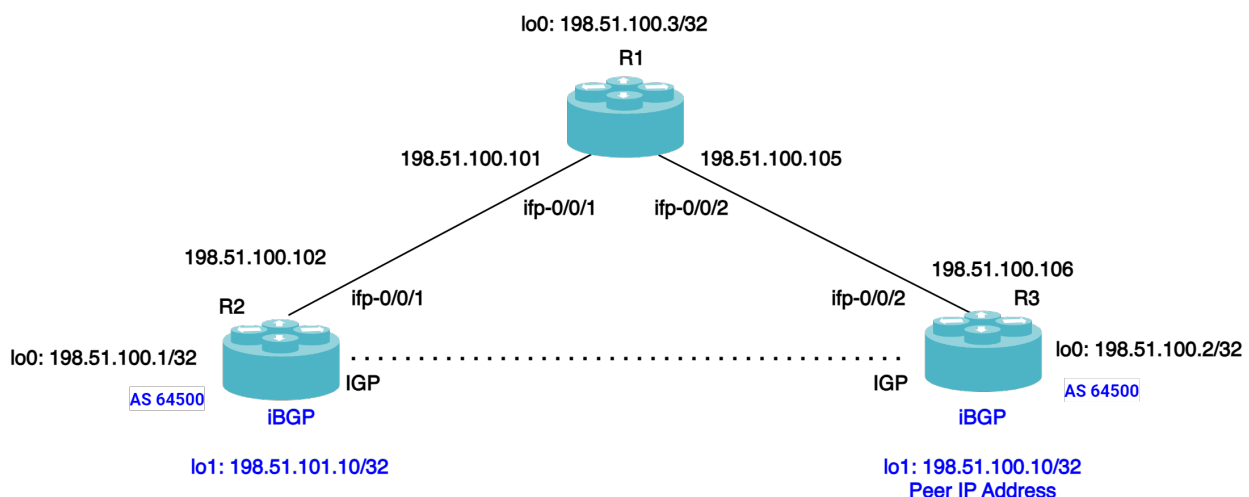
1.3. Route Selection

RIBD obtains routes and corresponding next-hops from various sources such as static, protocols, and direct connections. Route selection and resolution are also handled by RIBD, which chooses the best route based on the preference assigned to each source. If multiple sources provide the same prefix, RIBD will select the best route based on the source's preference and install it in FIBD.

1.4. Route Resolution

Route resolution is the process of finding the forwarding next hop from protocol nexthop downloaded into RIBD. For this purpose RIBD performs recursive route lookups till it finds a direct outgoing interface for the route.

For example, BGP installs a route with nexthop set to its peer IP address. However, it doesn't provide the outgoing interface. So, RIBD will check the peer IP address (next-hop IP) in its routing table (most probably downloaded by IGP) and find the directly connected router's IP address and its outgoing nexthop.



The BGP route for 198.51.100.10/32 will only be added to the routing table of router R2 if the IP address listed as the next-hop attribute is already reachable based on the information stored in the routing table. Additionally, the BGP route that is installed will contain a reference to the next-hop address 198.51.100.2 (refer to the figure above).

The network 198.51.100.10/32 can be accessed through an IP address that is not directly connected. The physical interface is not located where the BGP route is installed, so it is added to the IP routing table without any information about the outgoing interface. To find the BGP next-hop in the routing table, the router must perform a recursive lookup. However, to use a BGP route, the BGP next-hop IP address must be reachable. This is usually done through an IGP that provides reachability information. In this case, the BGP next-hop 198.51.100.2 is found in the routing table of R2, which is known via OSPF. The outgoing interface is ifp-0/0/1.

During the first route lookup, the router checks whether the destination prefix is in the routing table. If it is, then a recursive lookup is performed for the next-hop IP address. Since the next hop address is not a directly connected interface, the router needs to do a recursive lookup to find it.

Below is an example of the BGP route resolution for prefix:198.51.100.10/32:

```
supervisor@rtbrick>C-BNG.rtbrick.net: op> show route-resolution resolved prefix
198.51.100.10
198.51.100.10, Source: bgp
  Destination instance: default, AFI: ipv4, SAFI: unicast
  Lookup instance: default, AFI: ipv4, SAFI: unicast
  Covering Prefix: 198.51.100.10
    Interface      MAC Address      Nexthop
    if1-0/0/1/13   e8:c5:7a:8f:56:47  198.51.100.101
supervisor@rtbrick>C-BNG.rtbrick.net: op>
```

1.5. Policy Attachments

Routing Policies are a set of rules that enable you to manage and alter the default behavior of routing protocols, like BGP and IS-IS. Such policies consist of various "terms," which include "match" and "action" sections with control. The traffic that matches the "match" block is handled by the "action" block.

Once policies have been created, they need to be applied to take effect. Attachment points describe the specific applications and processes to which policies can be applied.

For more information, see section "2.2.4. Attaching Policies" of the Policy User Guide.

1.6. Nexthops

Nexthop helps routers determine the best path for forwarding data packets to their final destination efficiently. The next hop is identified by its IP address, which is stored in the routing table alongside the associated network prefix and other routing information.

1.7. Adjacency

The FIB learns the routing information from the routing table and tracks the next hop for all routes. The adjacency table maintains Layer 2 information(Nexthop) for the routes listed in the FIB (the resolved routes that can be installed into the hardware).

1.8. FIB

FIB determines how packets are forwarded (data plane), i.e., it only contains a subset of information from RIB. Also, FIB is downloaded to ASIC or VPP.

The main difference between the RIB and the FIB is that the RIB contains all the routes the router has learned, while the FIB only contains the best paths to each destination network. The RIB is constantly updated by the routing protocols and static configuration, while the FIB is only updated when the router needs to recalculate the best paths. The FIB is also updated more quickly than the RIB since it only needs to store the best paths and not all routes. The FIB maintains next-hop address information based on the information in the IP routing table.

2. Operational commands

2.1. Display Routes

The show route command displays information on routes.

Syntax:

show route <option>

Option	Description
-	Without any option, the command displays the information for all routes for all instances.
detail	Shows detailed route information.
instance <name>	Routing table information for a specified instance.
ipv4	Shows route information for the IPv4 routing table.
ipv6	Shows route information for the IPv6 routing table.
mpls	Shows route information for the MPLS routing table.
label <value>	Shows route information for a specified destination label.
prefix <value>	Shows route information for a specified destination prefix.
prefix-length-distribution	Shows the number of routes with the same prefix length for the sources.
source	Shows routes from a specified source.
summary	Shows the number of routes selected by RIBD for each source.

Example 1: The following example shows route information.

```

supervisor@rtbrick>LEAF01: op> show route
Instance: default, AFI: ipv4, SAFI: unicast
Prefix/Label          Source          Pref    Next Hop
Interface
192.1.0.3/32          direct         0      192.1.0.3
lo-0/0/0/1

Instance: default, AFI: ipv4, SAFI: labeled-unicast
Prefix/Label          Source          Pref    Next Hop
Interface             Label
192.1.0.3/32          direct         0      192.1.0.3
lo-0/0/0/1           -

Instance: default, AFI: ipv6, SAFI: unicast
Prefix/Label          Source          Pref    Next Hop
Interface
192:1::2/128         bgp            20     fe80::eac5:7aff:fe8f:5663
if1-0/1/70/13
192:1::3/128         direct         0      192:1::3
lo-0/0/0/1

192:1::4/128         bgp            20     fe80::eac5:7aff:fe8f:5663
if1-0/1/70/13
Instance: default, AFI: ipv6, SAFI: labeled-unicast
Prefix/Label          Source          Pref    Next Hop
Interface             Label
192:1::2/128         bgp            20     fe80::eac5:7aff:fe8f:5663
if1-0/1/70/13       -
192:1::3/128         direct         0      192:1::3
lo-0/0/0/1           -
192:1::4/128         bgp            20     fe80::eac5:7aff:fe8f:5663
if1-0/1/70/13       2020

```

Example 2: A route summary is shown in the following example.

```
supervisor@rtbrick>LEAF01: op> show route summary
Instance: default
  Source          Routes
  bgp              4
  direct           4
  Total Routes    8
Instance: inband-vrf
  Source          Routes
  bgp              4
  direct           2
  Total Routes    6
Instance: ip2vrf
  Source          Routes
  bgp              3
  direct           2
  Total Routes    5
Instance: li-vrf
  Source          Routes
  bgp              2
  direct           2
  Total Routes    4
Instance: mgmt-vrf
  Source          Routes
  bgp              4
  direct           2
  Total Routes    6
Instance: radius-vrf
  Source          Routes
  bgp              3
  direct           2
  Total Routes    5
```

Example 3: This example shows routes with IPv4 prefixes of the same length

```

supervisor@rtbrick>LEAF01: op>> show route prefix-length-distribution
Instance: default
  Prefix Length    Count
      /32          2
      /128         6
      Sum          8
Instance: inband-vrf
  Prefix Length    Count
      /32          3
      /128         3
      Sum          6
Instance: ip2vrf
  Prefix Length    Count
      /24          1
      /32          2
      /128         2
      Sum          5
Instance: li-vrf
  Prefix Length    Count
      /32          2
      /128         2
      Sum          4
Instance: mgmt-vrf
  Prefix Length    Count
      /32          3
      /128         3
      Sum          6
Instance: radius-vrf
  Prefix Length    Count
      /24          1
      /32          2
      /128         2
      Sum          5

```

2.2. Displaying Route Resolution

The `show route-resolution` command displays the routes that were requested to be resolved for their nexthops. Otherwise, it shows the route is unresolved.

Syntax:

show route-resolution <options>

Option	Description
-	Without any option, the command displays the information for all requests and response tables side by side.
destination-instance	Displays the information for all requests and responses for a destination instance.
look-up instance	Displays lookup instance routes.
prefix	Displays routes for prefix 4 or prefix 6.

Option	Description
resolved	Displays resolved routes.
source	Displays source of requested source.
unresolved	Displays unresolved routes.

Example 1: Below is an example of the route resolution requested by the protocol BGP.

```

supervisor@rtbrick>LEAF01: op> show route-resolution
192:1::2, Source: bgp
  Destination instance: default, AFI: ipv4, SAFI: vpn-multicast
  Lookup      instance: default, AFI: ipv6, SAFI: labeled-unicast
  Covering Prefix: 192:1::2/128
    Interface      MAC Address      Nexthop
    if1-0/1/70/13  e8:c5:7a:8f:56:63  fe80::eac5:7aff:fe8f:5663
fe80::eac5:7aff:fe8f:5663, Source: bgp
  Destination instance: default, AFI: ipv6, SAFI: unicast
  Lookup      instance: default, AFI: ipv6, SAFI: unicast
  Covering Prefix: fe80::eac5:7aff:fe8f:5663/128
    Interface      MAC Address      Nexthop
    if1-0/1/70/13  e8:c5:7a:8f:56:63  fe80::eac5:7aff:fe8f:5663
fe80::eac5:7aff:fe8f:5663, Source: bgp
  Destination instance: default, AFI: ipv4, SAFI: vpn-unicast
  Lookup      instance: default, AFI: ipv6, SAFI: labeled-unicast
  Covering Prefix: fe80::eac5:7aff:fe8f:5663/128
    Interface      MAC Address      Nexthop
    if1-0/1/70/13  e8:c5:7a:8f:56:63  fe80::eac5:7aff:fe8f:5663
fe80::eac5:7aff:fe8f:5663, Source: bgp
  Destination instance: default, AFI: ipv6, SAFI: labeled-unicast
  Lookup      instance: default, AFI: ipv6, SAFI: labeled-unicast
  Covering Prefix: fe80::eac5:7aff:fe8f:5663/128
    Interface      MAC Address      Nexthop
    if1-0/1/70/13  e8:c5:7a:8f:56:63  fe80::eac5:7aff:fe8f:5663
fe80::eac5:7aff:fe8f:5663, Source: bgp
  Destination instance: default, AFI: ipv6, SAFI: vpn-unicast
  Lookup      instance: default, AFI: ipv6, SAFI: labeled-unicast
  Covering Prefix: fe80::eac5:7aff:fe8f:5663/128
    Interface      MAC Address      Nexthop
    if1-0/1/70/13  e8:c5:7a:8f:56:63  fe80::eac5:7aff:fe8f:5663

```