

Chapter 6

Routing Protocols Concepts

The JUNOS routing protocol process (rpd) for IPv6 supports a variety of routing protocols, including Intermediate System to Intermediate System (IS-IS), Routing Information Protocol Next-Generation (RIPng), and Border Gateway Protocol (BGP). This chapter explains the general terminology and concepts related to configuring and using the routing protocol process and the routing protocols. For information about configuring the individual routing protocols, see the individual chapter about that protocol.

This chapter discusses the following topics:

Routing Databases on page 59

Route Preferences on page 62

Equal-Cost Paths and Load Sharing on page 65

Routing Databases

The JUNOS software maintains two databases for routing information:

Routing table—Contains all the routing information learned by all routing protocols.

Forwarding table—Contains the routes actually used to forward packets through the router.

In addition, link-state protocols such as IS-IS maintain link-state databases.

The following sections describe each database:

Routing Protocol Databases on page 60

JUNOS Routing Tables on page 60

Forwarding Tables on page 61

How the Routing and Forwarding Tables Are Synchronized on page 61

Routing Protocol Databases

Each interior gateway protocol (IGP) routing protocol maintains a database of routing information learned from other routers running the same protocol, and uses this information as defined and required by the protocol. IS-IS uses the routing information received to maintain link-state databases, which is used to determine which adjacent neighbors are operational and to construct network topology maps.

IS-IS uses the Dijkstra algorithm and RIPng uses the Bellman-Ford algorithm to determine the best route or routes (if there are multiple equal-cost routes) to reach each destination and installs these routes into the JUNOS software routing table.

JUNOS Routing Tables

The JUNOS software routing table is used by the routing protocol process to maintain its database of routing information. In this table, the routing protocol process stores statically configured routes, directly connected interfaces (also called *direct routes* or *interface routes*), and all routing information learned from all routing protocols. The routing protocol process uses this collected routing information to select the *active route* to each destination, which is the route that actually is used to forward packets to that destination.

By default, the JUNOS software maintains a routing table for unicast routes, another for multicast routes, and a third for Multiprotocol Label Switching (MPLS). You can configure additional routing tables to support situations where you need to separate a particular group of routes or where you need greater flexibility in manipulating routing information. In general, most operations can be performed without resorting to the complexity of additional routing tables. However, creating additional routing tables has several specific uses, including importing interface routes into more than one routing table, applying different routing policies when exporting the same route to different peers, and providing greater flexibility with incongruent multicast topologies.

Each routing table is identified by a name, which consists of the protocol family followed by a period and small, nonnegative integer. The protocol family can be inet6 (IPv6), inet (IPv4), or iso (ISO). The following names are reserved for the default routing tables maintained by the JUNOS software:

inet6.0—Default IPv6 unicast routing table

inet.0—Default IPv4 unicast routing table



Note

For clarity, this manual contains general discussions of routing tables as if there were only one table. However, when it is necessary to distinguish among the routing tables, their names are explicitly used.

Forwarding Tables

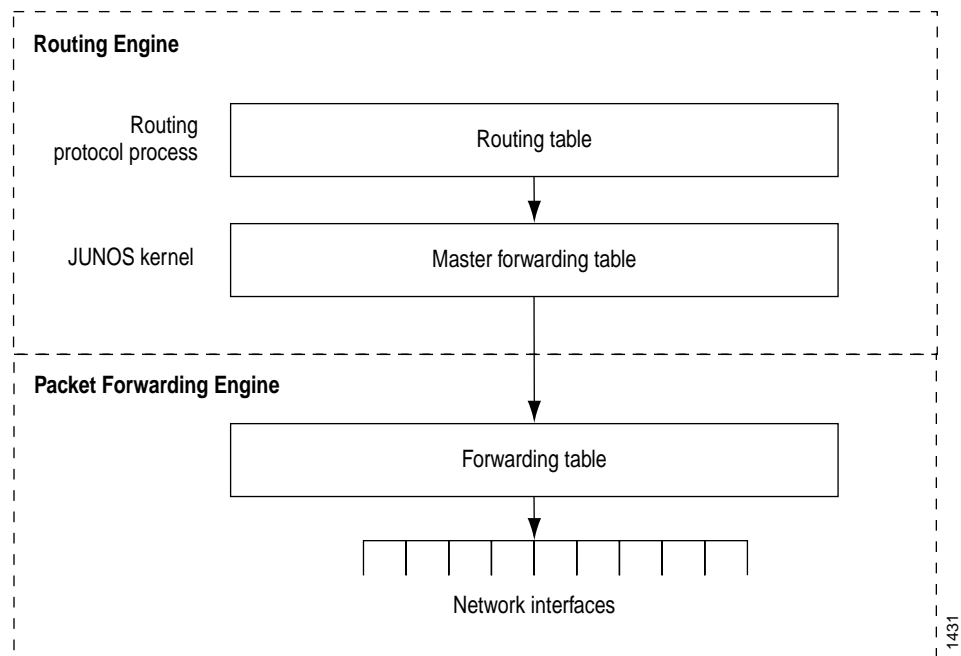
The JUNOS software installs all active routes from the routing table into the forwarding table. The active routes are routes that are used to forward packets to their destinations.

The JUNOS kernel maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router responsible for forwarding packets.

How the Routing and Forwarding Tables Are Synchronized

The JUNOS routing protocol process is responsible for synchronizing the routing information between the routing and forwarding tables. To do this, the routing protocol process calculates the active routes from all the routes in the routing table and installs them into the forwarding table. The routing protocol process then copies the forwarding table to the router's Packet Forwarding Engine, the part of the router that forwards packets. Figure 3 illustrates how the routing tables are synchronized.

Figure 3: Synchronizing Routing Exchange between the Routing and Forwarding Tables



Route Preferences

For unicast routes, the JUNOS routing protocol process uses the information in its routing table, along with the properties set in the configuration file, to choose an *active route* for each destination. While the JUNOS software might know of many routes to a destination, the active route is the preferred route to that destination and is the one that is installed in the forwarding table and used when actually routing packets.

The routing protocol process generally determines the active route by selecting the route with the lowest preference value. The preference is an arbitrary value in the range 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems.

When routes in the routing table are nearly identical, the routing protocol process prefers the route whose next hop has the lowest IP address.

The preference value is used to select routes to destinations in external autonomous systems (ASs) or routing domains; it has no effect on the selection of routes within an AS (that is, within an IGP). Routes within an AS are selected by the IGP and are based on that protocol's metric or cost value.

The following sections describe some of the algorithms used for handling route preferences:

Alternate and Tiebreaker Preferences on page 62

How the Active Route Is Determined on page 63

Multiple Active Routes on page 64

Default Route Preference Values on page 64

Alternate and Tiebreaker Preferences

The JUNOS software provides support for alternate and tiebreaker preferences, and some of the routing protocols, including BGP, use these additional preferences. With these protocols, you can specify a primary route preference, `preference`, and a secondary preference, `preference2`, that is used as a tiebreaker. You can also mark route preferences with additional route tiebreaker information by specifying a `color`, `color`, and a tiebreaker `color`, `color2`.

The software uses a four-byte value to represent the route preference value. When using the preference value to select an active route, the software first compares the primary route preference values, choosing the route with the lowest value. If there is a tie and if a secondary preference has been configured, the software compares the secondary preference values, choosing the route with the lowest value. If there is still a tie, the software continues the comparison process, comparing the configured primary and secondary color values.

How the Active Route Is Determined

For each prefix in the routing table, the routing protocol process selects a single best path, called the active route. The algorithm for determining the active route is as follows:

1. Choose the path with the lowest preference value (routing protocol process preference). Routes that are not eligible to be used for forwarding (for example, because they were rejected by routing policy or because a next hop is inaccessible) have a preference of -1 and are never chosen.
2. For BGP, prefer the path with higher local preference. For non-BGP paths, choose the path with the lowest preference2 value.
3. If the path includes an AS path:
 - a. Prefer the route with a shorter AS path.

Confederation sequences are considered to have a path length of 0, and AS and confederation sets are considered to have a path length of 1.

- b. Prefer the route with the lower origin code. Routes learned from an IGP have a lower origin code than those learned from an EGP, and both of these have lower origin codes than incomplete routes (routes whose origin is unknown).
- c. Depending on whether nondeterministic routing table path selection behavior is configured, there are two possible cases:

If nondeterministic routing table path selection behavior is not configured (that is, if the path-selection cisco-nondeterministic statement is not included in the BGP configuration), for paths with the same neighboring AS numbers at the front of the AS path, prefer the path with the lowest multiple exit discriminator (MED) metric. Confederation AS numbers are not considered when deciding what the neighbor AS number is. When you display the routes in the routing table using the show route command, they generally appear in order from most preferred to least preferred. Routes that share the same neighbor AS are grouped together in the command output. Within a group, the best route is listed first and the other routes are marked with the NotBest flag in the State field of the show route detail command.

To always compare MEDs whether or not the peer ASs of the compared routes are the same, use the path-selection (always-compare-med) statement. For an example, see “Configure Routing Table Path Selection” on page 230.

If nondeterministic routing table path selection behavior is configured (that is, the path-selection cisco-nondeterministic statement is included in the BGP configuration), prefer the path with the lowest MED metric. When you display the routes in the routing table using the show route command, they generally appear in order from most preferred to least preferred and are ordered with the best route first, followed by all other routes in order from newest to oldest.

In both cases, confederations are not considered when determining neighboring ASs. Also, in both cases, a missing metric is treated as if a MED were present but zero.

4. Prefer strictly internal paths, which include IGP routes and locally generated routes (static, direct, local, and so forth).
5. Prefer strictly external (EBGP) paths over external paths learned through interior sessions (IBGP).
6. For BGP, prefer the path whose next hop is resolved through the IGP route with the lowest metric.
7. Prefer paths from the higher routing table number. For example, inet.3 is preferred to inet.0.
8. Prefer paths with a larger number of next hops.
9. For BGP, prefer the route with the lowest IP address value for the BGP router ID.
10. Prefer the path that was learned from the neighbor with the lowest peer IP address.

Multiple Active Routes

The IGP's compute equal-cost multipath next hops, and IBGP picks up these next hops. When there are multiple, equal-cost next hops associated with a route, the routing protocol process installs only one of the next hops in the forwarding path with each route, randomly selecting which next hop to install. For example, if there are 3 equal-cost paths to an exit router and 900 routes leaving through that router, each path ends up with about 300 routes pointing at it. This mechanism provides load distribution among the paths while maintaining packet ordering per destination.

Default Route Preference Values

The JUNOS software routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference is a value from 0 through 255, with a lower value indicating a more preferred route. Table 2 lists the default preference values. In general, the narrower the scope of the statement, the higher precedence its preference value is given, but the smaller the set of routes it affects. To modify the default preference value for routes learned by routing protocols, you generally apply routing policy when configuring the individual routing protocols. You also can modify some preferences with other configuration statements, which are indicated in table 2.

Table 2: Default Route Preference Values

How Route Is Learned	Default Preference	Statement to Modify Default Preference
Directly connected network	0	—
Static	5	static on page 118
IS-IS Level 1 internal route	15	IS-IS external-preference on page 150, preference on page 158
IS-IS Level 2 internal route	18	IS-IS external-preference on page 150, preference on page 158
Redirects	30	—
RIPng	100	RIPng preference on page 180
Point-to-point interface	110	—
IS-IS Level 1 external route	160	IS-IS external-preference on page 150, preference on page 158
IS-IS Level 2 external route	165	IS-IS external-preference on page 150, preference on page 158
BGP	170	BGP preference on page 260, export on page 246, import on page 250

Equal-Cost Paths and Load Sharing

For equal-cost paths, load sharing is based on the BGP next hop. For example, if four prefixes all point to a next hop and there is more than one equal-cost path to that next hop, the routing protocol process randomizes the path chosen among the four prefixes. Also, for each prefix, the routing protocol process installs a single forwarding entry pointing along one of the paths. The routing software does not rerandomize the path taken as prefixes pointing to the next hop come and go, but it does rerandomize if the number of paths to the next hop changes. Because a prefix is tied to a particular path, packet reordering should not happen. The degree of load sharing improves as the number of prefixes increases.

