

Chapter 12

Configure Other Protocol-Independent Routing Properties

This chapter discusses how to perform the following tasks for configuring other protocol-independent routing properties:

- Configure the AS Number on page 121
- Configure the Router Identifier on page 122
- Configure AS Confederation Members on page 122
- Configure Route Recording for Flow Aggregation on page 123
- Create Routing Table Groups on page 123
- Configure How Interface Routes Are Imported into Routing Tables on page 124
- Configure Multicast Scoping on page 124
- Configure Per-Packet Load Balancing on page 125
- Configure Logging for the Routing Protocol Process on page 127
- Trace Global Routing Protocol Operations on page 128

Configure the AS Number

An AS is a set of routers that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routers. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

ASs are identified by a number from 1 through 65,535 that is assigned by the Network Information Center (NIC) (in the United States, <http://www.isi.edu>).

If you are using BGP on the router, you must configure an AS number.

To configure the router's AS number, include the `autonomous-system` statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  autonomous-system autonomous-system <loops number>;
}
```

To specify how many times this AS number can appear in an AS path, include the `loops` option.

Configure the Router Identifier

The router identifier is used by BGP and OSPF to identify the router from which a packet originated. The router identifier usually is the IP address of the local router. If you do not configure a router identifier, the IP address of the first interface encountered in the router is used.

To configure the router identifier, include the `router-id` statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  router-id address;
}
```

Configure AS Confederation Members

If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more *confederations*. Each confederation is identified by its own AS number, which is called a *confederation AS number*. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.

The BGP path attributes `NEXT_HOP`, `LOCAL_PREF`, and `MULTI_EXIT_DISC`, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.

Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.

Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.

If you are using BGP, you can enable the local router to participate as a member of an AS confederation. To do this, include the `confederation` statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  confederation confederation-autonomous-system members [autonomous-system];
}
```

Specify the AS confederation identifier, along with the AS numbers that are members of the confederation.

Note that peer adjacencies will not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.

Configure Route Recording for Flow Aggregation

Before you can perform flow aggregation, RPD must export the AS path and routing information to the sampling process. To do this, include the `route-record` statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  route-record;
}
```

For more information about flow aggregation and sampling, see the *JUNOS Internet Software Configuration Guide: Interfaces and Chassis*.

Create Routing Table Groups

You can group together one or more routing tables to form a *routing table group*. Within a group, a routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

To create a routing table group, include the `rib-groups` statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  rib-groups group-name {
    import-rib [ routing-table-name ];
    export-rib routing-table-name;
  }
}
```

The routing table group can have any name you choose (specified in *group-name*). If the group name you specify is not created explicitly, as described in “Configure Other Protocol-Independent Routing Properties” on page 121, naming it in the `rib-groups` statement creates it.

Each routing table group must contain one or more routing tables that the JUNOS software uses when importing routes (specified in the `import-rib` statement). The first routing table you specify is the *primary routing table*, and any additional routing tables are the *secondary routing tables*. The primary routing table must be `inet.0`.

Each routing table group optionally can contain one routing table group that the JUNOS software uses when exporting routes to the routing protocols (specified in the `export-rib` statement).

Example: Create a Routing Table Group

Create a routing table group so that interface routes are installed into two routing tables, inet.0 and inet.2:

```
[edit]
routing-options {
  interface-routes {
    rib-groups if-rg;
  }
  rib-groups if-rg {
    rib-group if-rg {
      import-rib [ inet.0 inet.2 ];
    }
  }
}
```

Configure How Interface Routes Are Imported into Routing Tables

By default, interface routes (also called direct routes) are imported into routing table inet.0 only. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the router's interfaces.

To associate a routing table group with the router's interfaces and specify which routing table groups interface routes are imported into, include the interface-routes statement at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  interface-routes {
    rib-group routing-table-name;
  }
}
```

To create the routing table groups, include the rib-groups statement at the [edit routing-options] hierarchy level. For configuration information, see "Create Routing Table Groups" on page 123.

Configure Multicast Scoping

To configure multicast address scoping, include the following statements at the [edit routing-options] hierarchy level:

```
[edit]
routing-options {
  multicast {
    scope scope-name {
      interface [ interface-name ];
      prefix destination-prefix;
    }
  }
}
```

Specify a name for the scope, its address range, and the router interfaces on which you are configuring scoping.

Example: Configure Multicast Scoping

Configure multicast scoping, creating four scopes, local, organization, engineering, and marketing:

```
[edit]
routing-options {
  multicast {
    scope local {
      prefix 239.255.0.0/16;
      interface ip-f/p/0;
    }
    scope organization {
      prefix 239.192.0.0/14;
      interface [ ip-f/p/0 so-0/0/0 ];
    }
    scope engineering {
      prefix 239.255.255.0/24;
      interface [ ip-f/p/0 so-0/0/1 so-0/0/2 ];
    }
    scope marketing {
      prefix 239.255.254.0/24;
      interface [ ip-f/p/0 so-0/0/3 so-1/0/0 ];
    }
  }
}
```

Configure Per-Packet Load Balancing

For the active route, when there are multiple equal-cost paths to the same destination, by default, the JUNOS software chooses in a random fashion one of the next-hop addresses to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen, also in a random fashion.

You can configure the JUNOS software so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routers. The behavior of per-packet load balancing function varies, according to the version of the Internet Protocol ASIC in the router.

On routers with an Internet Processor I ASIC, when per-packet load balancing is configured, traffic between routers with multiple paths is spread in a random fashion across the available interfaces. The forwarding table balances the traffic headed to a destination, transmitting it in round-robin fashion among the multiple next hops (up to a maximum of 8 equal-cost load-balanced paths). The traffic is load-balanced on a per-packet basis.

On routers with the Internet Processor II ASIC, when per-packet load balancing is configured, traffic between routers with multiple paths is divided into individual traffic flows (up to a maximum of 16 equal-cost load-balanced paths). Packets for each individual flow are kept on a single interface. To recognize individual flows in the transit traffic, the router examines each of the following:

- Source IP address
- Destination IP address
- Protocol

- Source port number

- Destination port number

- Interface through which the packet entered the router

The router recognizes packets that have all of these parameters identical, and it ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

To configure per-packet load balancing, follow these steps:

1. Define a load-balancing routing policy by including one or more policy-statement statements at the [edit policy-options] hierarchy level, defining an action of load-balance per-packet:

```
[edit]
policy-options {
  policy-statement policy-name {
    from {
      match-conditions;
      route-filter destination-prefix match-type <actions>;
      prefix-list name;
    }
    then {
      load-balance per-packet;
    }
  }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table—To do this, include the export statement at the [edit routing-options forwarding-table] hierarchy level:

```
[edit]
routing-options {
  forwarding-table {
    export policy-name;
  }
}
```

Examples: Configure Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
user@host# set policy-options policy-statement load-balancing-policy then load-balance
per-packet
[edit]
user@host# set routing-options forwarding-table export load-balancing-policy
[edit]
user@host# show
policy-options {
  policy-statement load-balancing-policy {
    then {
      load-balance per-packet;
    }
  }
}
```

```

routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}

```

Perform per-packet load balancing only for a limited set of routes:

```

[edit]
user@host# set policy-options policy-statement load-balancing-policy from
route-filter 192.168.10/24 orlonger;
[edit]
user@host# set policy-options policy-statement load-balancing-policy from
route-filter 9.114/16 orlonger;
[edit]
user@host# set policy-options policy-statement load-balancing-policy then load-balance
per-packet
[edit]
user@host# set routing-options forwarding-table export load-balancing-policy
[edit]
user@host# show
policy-options {
  policy-statement load-balancing-policy {
    from {
      route-filter 192.168.10/24 orlonger;
      route-filter 9.114/16 orlonger;
    }
    then {
      load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}

```

Configure Logging for the Routing Protocol Process

To control how much information the routing protocol process should log, include the options statement at the [edit routing-options] hierarchy level. Include the following form of the statement to log messages at one or more individual severity levels:

```

[edit]
routing-options {
  options syslog level;
}

```

Include the following form of the statement to log messages up to and including a particular severity level:

```

[edit]
routing-options {
  options syslog upto level;
}

```

Examples: Configure System Logging for the Routing Protocol Process

Configure the router to log messages of all severities:

```
[edit]
user@host# set routing-options options syslog upto emergency
[edit]
user@host# show
routing-options {
  options syslog upto emergency;
}
```

Configure the router to log only alert-level and critical-level messages:

```
[edit]
user@host# set routing-options options syslog alert critical
[edit]
user@host# show
routing-options {
  options syslog alert critical;
}
```

Trace Global Routing Protocol Operations

Global routing protocol tracing operations track all general routing operations and record them in a log file. Any global tracing operations that you configure are inherited by the individual routing protocols. To modify the global tracing operations for an individual protocol, configure tracing when configuring that protocol.

For a general discussion of tracing and of the precedence of multiple tracing operations, see tracing and logging operations in the *JUNOS Internet Software Configuration Guide: Installation and System Management*.

To trace the paths of multicast packets, use the `mtrace` command, as described in the *JUNOS Internet Software Command Reference*.

To configure global routing protocol tracing flags, include the `traceoptions` statement at the `[edit routing-options]` hierarchy level:

```
[edit routing-options]
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

You can specify the following global routing protocol tracing flags:

- all—Trace all tracing operations.
- general—Trace all normal operations and routing table changes (a combination of the normal and route trace operations).
- normal—Trace all normal operations.
- policy—Trace policy operations and actions.
- route—Trace routing table changes.
- state—Trace state transitions.
- task—Trace interface transactions and processing.
- timer—Trace timer usage.

You can specify the following tracing flag modifiers:

- detail—Provide detailed trace information.
- receive—Trace only packets being received.
- send—Trace only packets being transmitted.

The flags in a traceoptions flag statement are identifiers. When you use the set command to configure a flag, any flags that might already be set are not modified. In the following example, setting the csN tracing flag has no effect on the already configure detail flag. Use the delete command to delete a particular flag.

```
[edit protocols isis]
user@host# show
traceoptions {
  flag csN detail;
}
[edit protocols isis]
user@host# set traceoptions flag csN
[edit protocols isis]
user@host# show
traceoptions {
  flag csN detail;
}
user@host# delete traceoptions flag detail
[edit protocols isis]
user@host# show
traceoptions {
  flag csN;
}
```

Examples: Configure Global Tracing Operations

Log all globally traceable operations, saving the output in up to 10 files that are up to 10 MB in size:

```
[edit]
routing-options {
  traceoptions {
    file routing size 10m files 10;
    flag all;
  }
}
```

Log all unusual or abnormal traceable operations:

```
[edit]
routing-options {
  traceoptions {
    file routing size 10m files 10;
    flag all;
    flag normal disable;
  }
}
```

Log changes that occur in the JUNOS software routing table:

```
[edit]
routing-options {
  traceoptions {
    file routing size 10m files 10;
    flag route;
  }
}
```