

Chapter 14

Configure Ethernet Interfaces

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 Mbps over coaxial cables and later over unshielded twisted pairs using 10BaseT. More recently, 100BaseTX (Fast Ethernet, 100 Mbps) and Gigabit Ethernet (1 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet

- Gigabit Ethernet

- Management Ethernet interface, which is an out-of-band management interface within the router

- Internal Ethernet interface, which connects the Routing Engine to the Packet Forwarding Engine

This chapter discusses the following topics specific to configuring the different types of Ethernet interfaces in the router:

- Configure Fast Ethernet and Gigabit Ethernet Physical Interface Properties on page 138

- Configure 802.1Q VLANs on page 140

- Configure Static ARP Table Entries on page 141

- Configure VRRP on page 142

- Configure the Management Ethernet Interface on page 149

- Configure the Internal Ethernet Interface on page 149

- Example: Configure Fast Ethernet Interfaces on page 150

- Example: Configure Gigabit Ethernet Interfaces on page 150

Configure Fast Ethernet and Gigabit Ethernet Physical Interface Properties

To configure Fast Ethernet-specific physical interface properties, include the `fastether-options` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex)
fastether-options {
  (loopback | no-loopback);
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

To configure Gigabit Ethernet-specific physical interface properties, include the `gigether-options` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
gigether-options {
  (flow-control | no-flow-control);
  (loopback | no-loopback);
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

You can configure the following properties specific to Fast Ethernet and Gigabit Ethernet interfaces:

Configure MAC Address Filtering on page 138

Configure Loopback Mode on page 139

Configure Flow Control on page 139

Configure the Link Characteristics on page 140

Configure MAC Address Filtering

On Fast Ethernet and Gigabit Ethernet interfaces, you can enable source address filtering, which blocks all incoming packets to that interface. To enable the filtering, include the `source-filtering` statement at the `[edit interfaces interface-name fastether-options]` or `[edit interfaces interface-name gigether-options]` hierarchy level:

```
source-filtering;
```

When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the `source-address-filter` statement at the `[edit interfaces interface-name fastether-options]` or `[edit interfaces interface-name gigether-options]` hierarchy level:

```
source-address-filter {
  mac-address;
  <additional-mac-address>;
}
```

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where *n* is a decimal digit. To specify more than one address, include multiple *mac-address* options in the *source-address-filter* statement.

If the remote Ethernet card is changed, the interface will not be able to receive packets from the new card because it will have a different MAC address.

Configure Loopback Mode

By default, local Fast Ethernet or Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the loopback statement at the [edit interfaces *interface-name* *fastether-options*] or [edit interfaces *interface-name* *gigether-options*] hierarchy level:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the loopback statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the *no-loopback* statement at the [edit interfaces *interface-name* *fastether-options*] hierarchy level:

```
[edit interfaces interface-name fastether-options]
no-loopback;
```

Configure Flow Control

By default, the router imposes flow control to regulate the amount of traffic sent out a Gigabit Ethernet interface. This is useful if the remote side of the connection is a Gigabit Ethernet switch.

You can disable flow control if you want the router to permit unrestricted traffic. To disable flow control, include the *no-flow-control* statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level:

```
[edit interfaces interface-name gigether-options]
no-flow-control;
```

To explicitly reinstate flow control, include the *flow-control* statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level:

```
[edit interfaces interface-name gigether-options]
flow-control;
```

Configure the Link Characteristics

By default, the router's management Ethernet interface, fxp0, and any installed Fast Ethernet interfaces, autonegotiate whether to operate in full-duplex or half-duplex mode. All other interfaces can operate only in full-duplex mode. For Gigabit Ethernet, the link partner must also be set to full duplex.

To explicitly configure a Fast Ethernet or the management Ethernet to operate in either full-duplex or half-duplex mode, include the link-mode statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
link-mode (full-duplex | half-duplex);
```

Configure 802.1Q VLANs

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, the software supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or broadcast domain.

The software supports receiving and forwarding routed Ethernet frames with 802.1Q Virtual Local Area Network (VLAN) tags, and running VRRP over 802.1Q-tagged interfaces. To configure the router to receive and forward frames with 802.1Q VLAN tags, include the vlan-tagging statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
vlan-tagging;
```

Gigabit Ethernet interfaces can be partitioned; up to 1024 logical interfaces are allowed, one for each VLAN. To bind a VLAN ID to a logical interface, include the vlan-id statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]  
vlan-id number;
```



Note

Because IS-IS has an 8-bit limit for broadcast multi-access media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For further information on IS-IS capabilities, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure VLAN-CCC Encapsulation

Ethernet interfaces with VLAN tagging enabled can use VLAN Circuit Cross-Connect (CCC) encapsulation. To configure the encapsulation on a physical interface, include the encapsulation vlan-ccc statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
encapsulation vlan-ccc;
```

Ethernet interfaces in VLAN mode can have multiple logical interfaces, but in CCC mode VLAN IDs from 0 through 511 are reserved for normal VLANs and VLAN IDs from 512 through 1023 are reserved for CCC VLANs.

In general, you configure an interface's encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for some encapsulation types, including Ethernet VLAN-CCC, you also can configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the encapsulation statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-ccc;
```

You cannot configure a logical interface with an encapsulation of `vlan-ccc` unless you also configure the physical device with the same encapsulation. The logical interface must also have a VLAN ID in the range from 512 through 1023; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering.

Example: Configure VLAN-CCC Encapsulation

Configure VLAN-CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

Configure Static ARP Table Entries

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses. To configure static ARP table entries, include the `arp` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
arp ip-address mac mac-address <publish>;
```

The IP address that you specify must be part of the subnet defined in the enclosing address statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: `nnnn.nnnn.nnnn` or `nn:nn:nn:nn:nn:nn`. For example, 0011.2233.4455 or 00:11:22:33:44:55.

If you include the `publish` option, the router replies to ARP requests for the specified IP address.

The JUNOS software does not support proxy ARP.

Example: Configure Static ARP Table Entries

Configure two static ARP table entries on the router's management interface:

```

interfaces fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}

```

Configure VRRP

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP). VRRP allows hosts on a LAN to make use of redundant routers on that LAN without requiring anything more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in the following document:

RFC 2338, *Virtual Router Redundancy Protocol*

To configure VRRP, include the `vrrp-group` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level:

```

[edit interfaces interface-name unit logical-unit-number family inet address address]
vrrp-group group-number {
  virtual-address [addresses];
  priority number;
  advertise-interval seconds;
  authentication-type authentication;
  authentication-key key;
  (preempt | no-preempt);
  track {
    interface interface-name priority-cost cost;
  }
}
[edit protocols vrrp]
traceoptions {
  flag flag <flag-modifier> <disable>;
}

```

You can configure the following VRRP properties:

Configure Basic VRRP Support on page 143

Configure VRRP Authentication on page 144

- Configure the Advertisement Interval for the VRRP Master Router on page 145
- Configure a Backup Router to Preempt the Master Router on page 145
- Configure an Interface To Be Tracked on page 145
- Trace Operations on Interfaces on which VRRP is Enabled on page 146
- For a VRRP configuration example, see “Example: Configure VRRP” on page 147.

Configure Basic VRRP Support

To set up a basic VRRP configuration, configure VRRP groups on interfaces by including the following statements at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level. An interface can be a member of one or more VRRP groups.

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
vrp-group group-number {
  virtual-address [addresses];
  priority number;
}
```

For each group, you must configure the following:

Group number—Identifies the VRRP group. It can be a value from 0 through 255.

If you also enable MAC source address filtering on the interface as described in “Configure MAC Address Filtering” on page 138, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Addresses of one or more virtual routers that are members of the VRRP group—These are the virtual IP addresses associated with the virtual router in the VRRP group. Normally, you configure only one virtual IP address per group. The virtual IP addresses must be the same for all routers in the VRRP group.

In the addresses, specify the address only. Do not include a prefix length.

If you configure a virtual IP address to be the same as the interface’s address (the address configured with the address statement), the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255 and you must configure preemption by including the preempt statement. If you have multiple VRRP groups on an interface, the interface can be the master virtual router for only one of the groups.

If the virtual IP address you choose is not the same as the interface’s address, you must ensure that this address does not appear anywhere else in the router’s configuration. For example, check that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.

Priority for this router to become the master virtual router—Is a value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority for becoming the master router. The router with the highest priority within the group becomes the master router.

Within a single VRRP group, the master and backup routers cannot be the same router.

Configure VRRP Authentication

All VRRP protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, VRRP authentication is disabled. You can configure one of the following authentication methods. Each VRRP group must use the same method.

Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.

MD5 algorithm—Is used to create the authentication data field in the IP authentication header. This header is used to encapsulate the VRRP protocol data unit (PDU). The receiving router uses an authentication key (password) to verify the authenticity of the IP authentication header and VRRP PDU.

To enable authentication and specify an authentication method, include the `authentication-type` statement at the `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-number]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
authentication-type authentication;
```

authentication can be `none`, `simple`, or `md5`. The authentication type must be the same for all routers in the VRRP group.

If you included the `authentication-type` statement to select an authentication method, you can configure a key (password) on each interface by including the `authentication-key` statement at the `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-number]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
authentication-key key;
```

The key (password) is an ASCII string. For simple authentication, it can be 1 through 8 characters long. For MD-5 authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" "). The key must be the same for all routers in the VRRP group.

Configure the Advertisement Interval for the VRRP Master Router

By default, the master default virtual router sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master router is still operational. If the master router fails or becomes unreachable, the backup router with the highest priority value becomes the new master router.

To modify the time between the sending of VRRP advertisement packets, include the `advertise-interval` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
advertise-interval seconds;
```

The interval can range from 1 through 255 seconds. The interval must be the same for all routers in the VRRP group.

Configure a Backup Router to Preempt the Master Router

The router with the highest priority value is the master virtual default router. If a backup router that has a higher priority than the current master router is brought online, by default, that router becomes the master router. That is, the backup router preempts the current master router.

To prohibit the local router from preempting the master router even if the local priority value is greater than the master's priority, include the `no-preempt` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
no-preempt;
```

If the master router that owns the virtual IP address preempts, the backup router always relinquishes its master role regardless of its preempt mode value.

Configure an Interface To Be Tracked

VRRP can track whether an interface is up or down and dynamically change the priority of the VRRP group based on the state of the tracked interface, which might trigger a new master router election.

When interface tracking is enabled, you cannot configure a priority of 255, hereby designating the master router. For each VRRP group, 1 through 10 interfaces can be tracked.

To configure an interface to be tracked, include the `track` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group
group-number]
track {
  interface interface-name priority-cost cost;
}
```

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked interface is down, forcing a new master router election. The cost can range from 1 through 254. The sum of the costs for all tracked interfaces or routes must be less than or equal to the configured priority of the VRRP group.

Trace Operations on Interfaces on which VRRP is Enabled

To trace the operations of interfaces on which VRRP is enabled, include the `traceoptions` statement at the `[edit protocols vrrp]` hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  flag flag;
}
```

You can specify the following VRRP tracing flags:

`all`—Trace all VRRP operations.

`database`—Trace all database changes.

`general`—Trace all general events.

`interfaces`—Trace all interface changes.

`normal`—Trace all normal events.

`packets`—Trace all packets sent and received.

`state`—Trace all state transitions.

`timer`—Trace all timer events.

By default, VRRP logs the error, DCD configuration and routing socket events to the file `/var/log/vrrpd`.

Example: Configure VRRP

Configure one master (Router A) and one backup (Router B) virtual default router. Note that the address configured in the virtual-address statements differs from the addresses configured in the address statements.

```
Router A:
[edit]
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.20/24 {
          vrrp-group 27 {
            virtual-address 192.168.1.15;
            priority 254;
            authentication-type simple;
            authentication-key booJUM;
          }
        }
      }
    }
  }
}

Router B:
[edit]
interfaces {
  ge-4/2/0 {
    unit 0 {
      family inet {
        address 192.168.1.24/24 {
          vrrp-group 27 {
            virtual-address 192.168.1.15;
            priority 200;
            authentication-type simple;
            authentication-key booJUM;
          }
        }
      }
    }
  }
}
```

When configuring multiple VRRP groups on an interface, configure one to be the master virtual router for that group:

```
[edit]
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.20/24 {
          vrrp-group 2 {
            virtual-address 192.168.1.20;
            priority 255;
            advertise-interval 3;
            preempt;
          }
          vrrp-group 10 {
            virtual-address 192.168.1.55;
            priority 201;
            advertise-interval 3;
          }
          vrrp-group 1 {
            virtual-address 192.168.1.54;
            priority 22;
            advertise-interval 4;
          }
        }
      }
    }
  }
}
```

Configure VRRP and MAC source address filtering on a Gigabit Ethernet interface. The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```
[edit interfaces]
ge-5/2/0 {
  gigger-options {
    source-filtering;
    source-address-filter {
      00:00:5e:00:01:0a;<— Virtual MAC address
    }
  }
  unit 0 {
    family inet {
      address 192.168.1.10/24 {
        vrrp-group 10 { <— VRRP group number
          virtual-address 192.168.1.10;
          priority 255;
          preempt;
        }
      }
    }
  }
}
```

Configure the Management Ethernet Interface

The router's management Ethernet interface, fxp0, is an out-of-band management interface. You must configure an IP address and prefix length for this interface, which you commonly do when you first install the software:

```
[edit]
user@host# set interfaces fxp0 unit 0 family inet address address/prefix-length
[edit]
user@host# show
interfaces {
  fxp0 {
    unit 0 {
      family inet {
        address address/prefix-length;
      }
    }
  }
}
```



Caution

The management Ethernet interface must be configured for the router to function.

Configure the Internal Ethernet Interface

The internal Ethernet interface, fxp1, connects the Routing Engine with the System Control Board (SCB), System and Switch Board (SSB), Forwarding Engine Board (FEB), or Switching and Forwarding Module (SFM), depending on router model, in the Packet Forwarding Engine. The router software automatically configures this interface.



Caution

Do not modify or remove the configuration for the internal Ethernet interface that the software automatically configures. If you do, the router will stop functioning.

```
user@host> show configuration
...
interfaces {
  ...
  fxp1 {
    unit 0 {
      family tnp {
        address 1;
      }
    }
  }
  ...
}
...
```

Example: Configure Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use 802.3 encapsulation.

```
[edit]
user@host# set interfaces fe-fpc/pic/port unit 0 family inet address local-address
user@host# show
interfaces {
  fe-fpc/pic/port {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

Example: Configure Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces use 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-fpc/pic/port unit 0 family inet address local-address
user@host# show
interfaces {
  ge-fpc/pic/port {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

Note that the M160 two-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links. This PIC is only supported on the M160 Internet Backbone Router and it requires a Type 2 M160 FPC.

Each of the two interfaces on the PIC is named:

```
ge-fpc/pic/[0..1]
```

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.