Chapter 7

Configure Protocol Family and Address Interface Properties

For each logical interface, you must configure one or more protocol families and you can configure interface address properties. To do this, you can include the following statements at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
family family {
    filter {
        input filter-name;
        output filter-name;
        group filter-group-number;
    }
    mtu bytes;
    multicasts-only;
    no-redirects;
    primary;
    address address {
        arp ip-address mac mac-address <publish>;
        destination address;
        broadcast address;
        multipoint-destination destination-address (dlci dlci-identifier | vci vci-identifier);
        multipoint-destination destination-address {
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period seconds;
            shaping {
                (cbr rate | vbr peak rate sustained rate burst length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        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;
            }
        }
    }
}
```

This chapter describes the interface protocol and address properties that you can configure:

The following protocol family and address properties are described in other chapters:

Configure the Protocol Family

For each logical interface, you can configure one or more of the following protocols that run on the interface:

- inet—IP (Internet Protocol). You must configure this protocol family for the logical interface to support IP protocol traffic, including OSPF, BGP, and ICMP.
- iso—ISO. You must configure this protocol family for the logical interface to support IS-IS traffic.
- mpls—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- tnp—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the System Control Board (SCB), System and Switch Board (SSB), Forwarding Engine Board (FEB), or System and Forwarding Module (SFM), depending on router model, in the router's Packet Forwarding Engine. The JUNOS software automatically configures this protocol family on the router's internal interfaces only, as discussed in "Configure the Internal Ethernet Interface" on page 149.

To configure the logical interface's protocol family, include the family statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level, specifying the desired family. To configure more than one protocol family on a logical interface, include multiple family statements.

```
[edit interfaces interface-name unit logical-unit-number]
family family {
    mtu size;
    multicasts-only;
    no-redirects;
    primary;
    address address {
        destination address;
        broadcast address;
        primary;
        preferred;
    }
}
```

Configure the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the inet family, you configure the interface's IP address. For the iso family, you configure an address for the loopback interface only. You can configure only one ISO address per interface. For the mpls and tnp families, you never configure an address.

To assign an address to an interface, include the address statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
address address {
    destination address;
    broadcast address;
    primary;
    preferred;
}
```

In the address statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the destination statement.

Broadcast address for the interface's subnet—Specify this in the broadcast statement.

Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet. For more information about preferred addresses, see "Configure Default, Primary, and Preferred Addresses and Interfaces" on page 83.

By default, the preferred address is the lowest numbered address on the subnet. To override the default and explicitly configure the preferred address, include the preferred statement when configuring the address.

Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you originate packets out the interface where the destination gives no hint about the subnet (for example, some ping commands). For more information about primary addresses, see "Configure Default, Primary, and Preferred Addresses and Interfaces" on page 83.

By default, the primary address on an interface is the lowest numbered non-127 preferred address on the interface. To override the default and explicitly configure the preferred address, include the primary statement when configuring the address.

Configure an Unnumbered Interface

When you need to conserve IP addresses, you can configure unnumbered interfaces. To do this, configure the protocol family, but do not include the address statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family*] hierarchy level:

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

For example:

```
[edit]
interfaces {
    so-6/1/0 {
        unit 0 {
            family inet;
            family iso;
        }
    }
}
```

When configuring unnumbered interfaces, you must ensure that a source address is configured on some interface in the router. This address is the default address. We recommend that you do this by assigning an address to the loopback interface (lo0), as described in "Configure the Loopback Interface" on page 157. If you configure an address (other than a martian) on the lo0 interface, that address is always the default address, which is desirable because the loopback interface is independent of any physical interfaces and therefore is always reachable.

Set the Protocol MTU

For each interface, you can configure an interface-specific MTU by including the mtu statement at the [edit interfaces interface *interface-name*] hierarchy level. If you need to modify this MTU for a particular protocol family, include the mtu statement at the [edit interfaces interface *interface-name* unit *logical-unit-number* family *family*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
mtu mtu;
```

The default protocol MTU is 4470 bytes for ATM PVC, Cisco HDLC, Frame Relay, and PPP encapsulations. For Ethernet encapsulation on IPv4, the default protocol MTU is 1500 bytes. For Ethernet encapsulation on ISO, the default protocol MTU is 1497 bytes.

> **Note**
>
> When you initially configure an interface, the protocol MTU is calculated automatically. However, if you subsequently change the media MTU, the protocol MTU on existing address families does not automatically adjust.
>
> If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the mtu statement at the [edit interfaces *interface-name*] hierarchy level, as discussed in "Configure the Media MTU" on page 30.)

When the family is mpls, the default protocol MTU is 1488 bytes. MPLS packets are 1500 bytes and have 4 to 12 bytes of overhead.The maximum number of DCLIs is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

Disable the Sending of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the no-redirects statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
no-redirects;
```

To disable the sending of protocol redirect messages for the entire router, include the no-redirects statement at the [edit system] hierarchy level.

Configure Default, Primary, and Preferred Addresses and Interfaces

The router has a default address and a primary interface, and interfaces have primary and preferred addresses.

The *default address* of the router is used as the source address on unnumbered interfaces. The routing protocol process tries to pick the default address as the router ID, which is used by protocols, including OSPF and IBGP.

The *primary interface* for the router is the interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. An interface's *preferred address* is the default local address used for packets sourced by the local router to destinations on the subnet.

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface lo0 that is not 127.0.0.1 is used.
2. The primary address on the primary interface is used.

To configure these addresses and interfaces, you can do the following:

Configure the Primary Interface for the Router on page 84

Configure the Primary Address for an Interface on page 84

Configure the Preferred Address for an Interface on page 85

Configure the Primary Interface for the Router

The *primary interface for the router* is the interface that packets go out when you type a command such as ping 255.255.255.255 or ping 224.0.0.1—that is, a command that does not include an interface name (that is, there is no interface *type*-0/0/0.0 qualifier) and the destination address does not imply any particular outgoing interface. It also is the interface on which multicast applications running locally on the router, such as SAP, do group joins by default. Finally, the primary interface also is the interface from which the default local address is derived for packets sourced out an unnumbered interface in the case where there are no non-127 addresses configured on the loopback interface, lo0.

By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest index address is chosen. Otherwise, any interface with an address could be picked. In practice, this means that, on the router, the fxp0 interface is picked by default.

To choose a different interface as the primary interface, include the primary statement when configuring the interface family:

```
[edit interfaces interface-name unit logical-unit-number family family]
primary;
```

Configure the Primary Address for an Interface

The *primary address on an interface* is the address that is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. For example, the local address in the packets sent by a ping interface so-0/0/0.0 255.255.255.255 command is the primary address on interface so-0/0/0.0. The primary address flag also can be useful for selection the local address used for packets sourced out unnumbered interfaces when multiple non-127 addresses are configured on the loopback interface, lo0. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

To set a different primary address, include the primary statement when configuring the subnet address:

```
[edit interfaces interface-name unit logical-unit-number family family address address]
primary;
```

Configure the Preferred Address for an Interface

The *preferred address on an interface* is the default local address used for packets sourced by the local router to destinations on the subnet. By default, the numerically lowest local address is chosen. For example, if the addresses 128.100.1.1/24, 128.100.1.2/24, and 128.100.1.3/24 are configured on the same interface, the preferred address on the subnet (by default, 128.100.1.1) would be used as a local address when you issue a ping 128.100.1.5 command.

To set a different preferred address for the subnet, include the preferred statement when configuring the subnet address:

```
[edit interfaces interface-name unit logical-unit-number family family address address]
preferred;
```

Configure a Point-to-Multipoint ATM Connection

To configure a point-to-multipoint (NBMA) ATM connection, include the multipoint-destination 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]
multipoint-destination destination-address vci vpi-identifier.vci-identifier;
```

You can also include the multipoint-destination 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]
multipoint-destination destination-address {
  inverse-arp;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period seconds;
  shaping {
    (cbr rate | vbr peak rate sustained rate burst length);
    queue-length number;
  }
  vci vpi-identifier.vci-identifier;
}
```

address is the interface's address. The address must include the destination prefix (for example, /24).

For each destination, include one multipoint-destination statement. *destination-address* is the address of the remote side of the connection, and *vci-identifier* and *vpi-identifier* are the VCI and optional VPI identifiers for the connection.

When you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

When configuring point-to-multipoint ATM connections, you can do the following:

Configure ATM Inverse ARP on page 86

Define the ATM Traffic-Shaping Profile on page 86

Define the ATM OAM F5 Loopback Cell Period on page 90

Configure the ATM OAM F5 Loopback Cell Threshold on page 90

Configure ATM Inverse ARP

You can configure ATM interfaces to support inverse ATM ARP, as described in RFC 2225. When inverse ATM ARP in enabled, the router responds to received Inverse ATM ARP requests by providing IP address information to the requesting ATM device.

The router does not initiate inverse ATM ARP requests.

By default, inverse ATM ARP is disabled. To configure a VC to respond to inverse ATM ARP requests, include the inverse-arp statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
address address {
    inverse-arp;
}
```

Define the ATM Traffic-Shaping Profile

When you are using an ATM encapsulation, you can configure a traffic-shaping profile that defines the following:

Bandwidth utilization, which consists of either a constant rate, or a peak cell rate with sustained cell rate and burst tolerance

Maximum queue length

These values are used in the ATM generic cell-rate algorithm, which is a leaky bucket algorithm that defines the short-term burst rate for ATM cells, the maximum number of cells that can be included in a burst, and the long-term sustained ATM cell traffic rate. Each individual VC has its own independent shaping parameters.

By default, the bandwidth utilization is unlimited; that is, unspecified bit rate (UBR) is used. Also, by default, buffer usage by VCs is unregulated. To define limits to bandwidth utilization on a point-to-point interface or to limit buffer use, include the shaping statement. For point-to-point interfaces, include the shaping statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name]
unit logical-unit-number {
    vci vpi-identifier.vci-identifier;
    shaping {
        (cbr rate | vbr peak rate sustained rate burst length);
        queue-length number;
    }
}
```

For virtual circuits that are part of a point-to-multipoint interface, include the shaping statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
address address {
    multipoint-destination destination-address {
        vci vpi-identifier.vci-identifier;
        shaping {
            (cbr rate | vbr peak rate sustained rate burst length);
            queue-length number;
        }
    }
}
```

Configure CBR

For traffic that does not require the ability to periodically burst to a higher rate, you can configure a constant bit rate (CBR) by including the cbr statement at the [edit interfaces *interface-name* unit *logical-unit-number*] or [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
    shaping {
        cbr rate;
    }
}
```

Configure VBR

To define variable bandwidth rate (VBR) utilization, include the vbr statement at the [edit interfaces *interface-name* unit *logical-unit-number*] or [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
vbr peak rate sustained rate burst length;
```

You can define the following VBR traffic-shaping properties:

Peak rate—Top rate at which traffic can burst.

Sustained rate—Normal traffic rate averaged over time.

Burst length—Maximum number of cells that a burst of traffic can contain. It can be a value from 1 through 255 cells.

Specify the rates in bits per second (bps) or cells per second (cps). For OC-3 interfaces, the highest rate is 135631698 bps (353207.55 cps), which corresponds to 100 percent of the available line rate. For OC-12 interfaces, the highest rate is 271263396 bps (706415.09 cps), which corresponds to 50 percent of the available line rate. Table 6 lists some of the other rates you can specify. If you specify a rate that is not listed, it is rounded to the nearest rate.

In general, the actual packet rate on the interface is calculated with the following formula:

```
actual-rate = (128 * line-rate) / (trunc ((128 * line-rate) / desired-rate))
```

line-rate is the maximum available rate on the interface (in bits per second) after factoring out the overhead for SONET and ATM (per-cell) overheads. For OC-3 interfaces, the line rate is calculated as follows:

line-rate = 155,520,000 bps x (26/27) x (48/53) = 135,631,698.1 bps

For OC-12 interfaces, the line rate is calculated as follows:

line-rate = 622,080,000 bps x (26/27) x (48/53) = 542,526,792.45 bps

desired-rate is the rate you enter in the vbr statement, in bits per second.

The trunc operator indicates that all digits to the right of the decimal point should be dropped.

For OC-3 interfaces, the maximum available rate is 100 percent of *line-rate*, or 135,631,698 bps. For OC-12 interfaces, the maximum available rate is 50 percent of *line-rate*, or 271,263,396 bps.

The following example shows the calculations for determining the actual rate when the desired rate is 80 percent of the maximum rate:

OC-3:
135,631,968 bps * 0.8 = 108,505,358 bps
actual-rate = (128 * 135,631,698.1) / (trunc ((128 * 135,631,698.1) / 101,723,773.5))
actual-rate = 17,360,857,344 / (trunc ((17,360,857,344) / 101,723,773.5)
actual-rate = 17,360,857,344 / 160
actual-rate = 108,505,358 bps

OC-12:
271,263,396 bps * 0.8 = 217,010,716.8 bps
actual-rate = (128 * 542,526,792.45) / (trunc ((128 * 542,526,792.45)/217,010,716.8))
actual-rate = 69,443,429,434 / (trunc ((69,443,429,434) / 217,010,716.8))
actual-rate = 69,443,429,434 / 320
actual-rate = 217,010,717 bps

Table 6: Traffic-Shaping Rates

Interface Type	Line Rate (bps)	Line Rate (cps)	Percentage of Total Line Rate
OC-3			
	135631698	353207.55	100.00
	134580290	350469.50	99.22
	133545057	347773.58	98.46
	132525629	345118.82	97.71
	131521647	342504.29	96.97
	130532762	339929.07	96.24
	129558637	337392.28	95.52
	128598943	334893.08	94.81
	127653363	332430.63	94.12
	126721587	330004.13	93.43
OC-12			
	271263396	706415.09	50.00
	270207897	703666.40	49.81
	269160579	700939.01	49.61
	268121349	698232.68	49.42
	267090113	695547.17	49.23
	266066779	692882.24	49.04
	265051257	690237.65	48.85
	264043458	687613.17	48.67
	263043293	685008.58	48.48
	262050677	682423.64	48.30

Buffers are shared among all VCs, and by default, there is no limit to the buffer size for a VC. If a VC is particularly slow, it might use all the buffer resources. To limit the queue size of a particular VC, include the queue-length statement when configuring the VC at the [edit interfaces *interface-name* unit *logical-unit-number*] or [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
queue-length number;
```

The length can range from 1 through 16383 packets. The default is 16383 packets.

Define the ATM OAM F5 Loopback Cell Period

When you are using an ATM encapsulation, you can configure the OAM F5 loopback cell period on virtual circuits, which is the interval at which OAM F5 loopback cells are transmitted.

By default, no OAM F5 loopback cells are sent. To send OAM F5 loopback cells on a point-to-point interface, include the oam-period statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name]
unit logical-unit-number {
    vci vpi-identifier.vci-identifier;
    oam-period seconds;
}
```

To send OAM F5 loopback cells on a virtual circuit that is part of a point-to-multipoint interface, include the oam-period statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
address address {
    multipoint-destination destination-address {
        vci vpi-identifier.vci-identifier;
        oam-period seconds;
    }
}
```

The period can range from 1 through 900 seconds.

OAM VC-AIS (alarm indication signal) and VC-RDI (remote defect indication) defect indication cells are used for identifying and reporting VC defects end-to-end. When a physical link or interface failure occurs, intermediate nodes insert OAM AIS cells into all the downstream VCs affected by the failure. Upon receiving an AIS cell on a VC, the router marks the logical interface down and sends an RDI cell on the same VC to let the remote end know the error status. When an RDI cell is received on a VC, the router sets the logical interface status to down. When no AIS or RDI cells are received for 3 seconds, the router sets the logical interface status to up. You do not need to configure anything to enable defect indication.

Configure the ATM OAM F5 Loopback Cell Threshold

When you are using an ATM encapsulation, you can configure the OAM F5 loopback cell threshold on VCs, which is the minimum number of consecutive OAM F5 loopback cells received before declaring that a VC is up or lost before declaring that a VC is down.

By default, when five consecutive OAM F5 loopback cells are received, the VC is considered to be up, and when five consecutive cells are lost, the VC is considered to be down. To modify these values on a point-to-point interface, include the oam-liveness statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name]
unit logical-unit-number {
    vci vpi-identifier.vci-identifier;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
```

```
}
```

To modify the OAM F5 loopback cell count threshold on a virtual circuit that is part of a point-to-multipoint interface, include the oam-liveness statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
address address {
    multipoint-destination destination-address {
        vci vpi-identifier.vci-identifier;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
    }
}
```

The cell count can be a value from 1 through 255 cells.

Configure a Point-to-Multipoint Frame Relay Connection

To configure a point-to-multipoint Frame Relay connection (also called a multipoint NBMA connection), include the multipoint-destination statement within the address statement at the [edit interfaces *interface-name* unit *logical-unit-number* address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
multipoint-destination destination-address dlci dlci-identifier;
```

address is the interface's address.

For each destination, include one multipoint-destination statement. *destination-address* is the address of the remote side of the connection, and *dlci-identifier* is the DLCI identifier for the connection.

When you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

If keepalives are enabled, causing the interface to send LMI messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. See the section "Configure Keepalives" on page 34.

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 93

Configure VRRP Authentication on page 94

Configure the Advertisement Interval for the VRRP Master Router on page 95

Configure a Backup Router to Preempt the Master Router on page 95

Configure an Interface To Be Tracked on page 95

Trace Operations on Interfaces on which VRRP is Enabled on page 96

For a VRRP configuration example, see "Example: Configure VRRP" on page 97.

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]
vrrp-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 51, 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—This value is 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 to 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 that 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 routers. 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 {
    gigether-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;
                }
            }
        }
    }
}
```

Apply Firewall Filters

To apply firewall filters to an interface, include the filter statements when configuring the logical interface at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] hierarchy level:

```
[edit interfaces]
interfaces interface-name {
    unit logical-unit-number {
        family inet {
            filter {
                group group-number;
                input filter-name;
                output filter-name;
            }
        }
    }
}
```

In the group statement, specify the interface group number to associate with the filter.

In the input statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the output statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.

You can use the same filter one or more times.

If you apply the filter to the interface lo0, it is applied to packets received or transmitted by the Routing Engine.

For more information about firewall filters, see "Firewall Filter Configuration Guidelines" on page 267.

Define Interface Groups in Firewall Filters

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You can then match these packets using the interface-group match statement, as described in Table 11 on page 275.

To define the interface to be part of an interface group, include the group statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet filter] hierarchy level:

```
[edit interfaces]
interfaces interface-name {
    unit logical-unit-number {
        family inet {
            filter {
                group group-number;
            }
        }
    }
}
```