

Chapter 32

Router Chassis Configuration Guidelines

You can configure properties of the router chassis, including the clock source, conditions that activate the red and yellow alarm light-emitting diodes (LEDs) on the router's craft interface, and SONET/SDH framing and concatenation properties for individual Physical Interface Cards (PICs).

To configure router chassis properties, you include statements at the [edit chassis] hierarchy level of the configuration:

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    sonet {
      device-count number;
    }
  }
  alarm {
    interface-type {
      alarm-name (red | yellow | ignore);
    }
  }
  fpc slot-number {
    pic pic-number {
      atm-cell-relay-accumulation;
      atm-l2circuit-mode (cell | aal5);
      vtmapping number;
      ce1 {
        e1 port-number {
          channel-group group-number timeslots slot-number;
        }
      }
      ct3 {
        port port-number {
          t1 link-number {
            channel-group group-number timeslots slot-number;
          }
        }
      }
    }
    framing (sdh | sonet);
    idle-cell-format {
      itu-t;
      payload-pattern;
    }
  }
}
```

```

mlfr-uni-nni-bundles number;
no-concatenate;
sparse-dlcis;
vtmapping (itu-t | klm);
}
}
(packet-scheduling | no-packet-scheduling);
vrf-mtu-check;
(source-route | no-source-route);
redundancy {
failover on-loss-of-keepalives;
graceful-switchover (disable | enable);
keepalive-time seconds;
routing-engine slot-number (master | backup | disabled);
sfm slot-number (always | preferred);
ssb slot-number (always | preferred);
}
}

```

This chapter describes the following tasks for configuring the router chassis:

- Minimum Chassis Configuration on page 485
- Configure Aggregated Devices on page 485
- Configure ATM Cell-Relay Accumulation Mode on page 485
- Configure Conditions That Trigger Alarms on page 486
- Configure SONET/SDH Framing on page 488
- Configure Sparse DLCIS Mode on page 489
- Configure Channelized PIC Operation on page 489
- Channelized DS-3 to DS-0 Naming on page 490
- Channelized E1 Naming on page 492
- Channelized STM-1 Interface Virtual Tributary Mapping on page 493
- Configure the Drop Policy for Traffic with Source-Route Constraints on page 494
- Configure Layer 2 Circuit Cell-Relay or AAL5 Transport Mode on page 494
- Configure Redundancy on page 494
- Configure Packet Scheduling on page 502
- Configure the Link Services PICs on page 502
- Configure Layer 2 Circuit Cell-Relay or AAL5 Transport Mode on page 494

Minimum Chassis Configuration

All the statements at the [edit chassis] hierarchy level of the configuration are optional.

Configure Aggregated Devices

JUNOS software supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard. To define the virtual links, you need to specify the associations between physical and logical devices within the [edit interfaces] hierarchy, and assign the correct number of logical devices by including the device-count statement at the [edit chassis aggregated-devices ethernet] and [edit chassis aggregated-devices sonet] hierarchy levels:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count 16;
  }
  sonet {
    device-count 16;
  }
}
```

The maximum number of logical devices you can assign is 16. For more information on physical and logical interfaces using aggregated links, including sample configurations, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Configure ATM Cell-Relay Accumulation Mode

You can configure ATM 1 PICs to use cell-relay accumulation mode. In this mode, the incoming cells (1 to 8 cells) are packaged into a single packet and forwarded to the label-switched path (LSP). At the edge router, this packet is divided into individual cells and transmitted over the ATM interface.



Note

When you configure an ATM PIC to use cell-relay accumulation, all ports on the ATM PIC use cell-relay accumulation mode.

To configure an ATM PIC to use cell-relay accumulation mode, include the atm-cell-relay-accumulation statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ]
atm-cell-relay-accumulation;
```

Configure Conditions That Trigger Alarms

For the different types of PICs, you can configure which conditions trigger alarms and whether they trigger a red or yellow alarm, or are ignored. Red alarm conditions light the RED ALARM LED on the router's craft interface and trigger an audible alarm if one is connected to the contacts on the craft interface. Yellow alarm conditions light the YELLOW ALARM LED on the router's craft interface and trigger an audible alarm if one is connected to the craft interface.

To configure conditions that trigger alarms and that can occur on any interface of the specified type, include the alarm statement at the [edit chassis] hierarchy level:

```
[edit chassis]
alarm {
  interface-type {
    alarm-name (red | yellow | ignore);
  }
}
```

alarm-name is the name of an alarm. Table 19 lists the systemwide alarms and the alarms for each interface type.

Table 19: Configurable PIC Alarm Conditions

Interface/System	Alarm Condition	Configuration Option
SONET/SDH and ATM	Link alarm indication signal	ais-l
	Path alarm indication signal	ais-p
	Signal degrade (SD)	ber-sd
	Signal fail (SF)	ber-sf
	Loss of cell delineation (ATM only)	locd
	Loss of framing	lof
	Loss of light	lol
	Loss of pointer	lop-p
	Loss of signal	los
	Phase locked loop out of lock	pll
	STS payload label (C2) mismatch	plm-p
	Line remote failure indication	rfl-l
	Path remote failure indication	rfl-p
	STS path (C2) unequipped	uneq-p

Interface/System	Alarm Condition	Configuration Option
E3/T3	Alarm indicator signal	ais
	Excessive numbers of zeros	exz
	Failure of the far end	ferf
	Idle alarm	idle
	Line code violation	lcv
	Loss of frame	lof
	Loss of signal	los
	Phase locked loop out of lock	pll
	Yellow alarm	ylw
Ethernet	Link has gone down	link-down
DS-1	Alarm indicator signal	ais
	Yellow alarm	ylw
Management-Ethernet	Link has gone down	link-down

Chassis Conditions That Trigger Alarms

Various conditions related to the chassis components trigger yellow and red alarms. You cannot configure these conditions. Table 20 lists the alarms that the chassis components can generate.

Table 20: Chassis Component Alarm Conditions

Chassis Component	Alarm Condition	Alarm Severity
Alternative Media	Router boots from alternate boot device. For more information about alternate boot devices, see "Boot Devices" on page 98.	Yellow
Fans	One fan has been removed from the chassis.	Yellow
	Two or more fans have been removed from the chassis.	Red
	One fan in the chassis is installed but not spinning.	Red
Power supplies	A power supply has been removed from the chassis.	Yellow
	A power supply has failed. If both power supplies fail, the router shuts down and the software might report the failures in the syslog file.	Red
Temperature	Chassis temperature has exceeded 54 degrees Centigrade and the fans have been turned on to full speed.	Yellow
	Chassis temperature has exceeded 75 degrees Centigrade and the router has been shut down.	Red
	The temperature sensor has failed.	Red
SCB/SSB/FEB/SFM	The control board (SCB, SSB, FEB, or SFM, depending on model) has failed. If this occurs, the board attempts to reboot.	Red

Chassis Component	Alarm Condition	Alarm Severity
FPC	An FPC has failed. If this occurs, the FPC attempts to reboot. If the SCB sees that an FPC is rebooting too often, it shuts down the FPC.	Red
Craft interface	The craft interface has failed.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's backplane from the front (generally, an FPC) is broken.	Red

Silence External Devices

You can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button located on the craft interface front panel. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after an external device is silenced reactivate the external device.

Configure SONET/SDH Framing

By default, SONET/SDH PICs use SONET framing. For a discussion of the differences between the two standards, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*. To configure a PIC to use SDH framing, include the framing statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level, specifying the sdh option:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number framing sdh
[edit chassis]
user@host# show
fpc slot-number {
  pic pic-number {
    framing sdh;
  }
}
```

To explicitly configure a PIC to use SONET framing, include the framing statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level, specifying the sonet option:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number framing sonet
[edit chassis]
user@host# show
fpc slot-number {
  pic pic-number
    framing sonet;
  }
}
```

Configure Sparse DLCIS Mode

By default, original Channelized DS-3 and original Channelized STM1 to E1 (or T1) interfaces can support a maximum of 64 data-link connection identifiers (DLCIs) per channel—as many as 1,792 DLCIs per DS-3 interface or 4,032 DLCIs per STM1 interface (0 through 63).

In sparse DLCIS mode, the full DLCI range (1 through 1,022) is supported. This allows you to use circuit cross-connect (CCC) and translation cross connect (TCC) features by means of Frame Relay on T1 and E1 interfaces. For more information about CCC and DLCIs, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.



Note

Sparse DLCIS mode requires a Channelized STM1 or Channelized DS-3 PIC.

DLCI 0 is reserved for LMI signaling.

Channelized T3 QPP and STM-1 QPP interfaces support a maximum of 64 DLCIs, numbered 0 through 1,022, and, therefore, do not require sparse mode.

To configure the router to use sparse DLCIS mode, include the `sparse-dlcis` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ]
sparse-dlcis;
```

Configure Channelized PIC Operation

By default, SONET PICs (interfaces with names `so-fpc/pic/port`) operate in concatenated mode, a mode in which the bandwidth of the interface is in a single channel. To configure a PIC to operate in channelized (multiplexed) mode, include the `no-concatenate` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number no-concatenate
[edit chassis]
user@host# show
fpc slot-number {
  pic pic-number
  no-concatenate;
}
```

When configuring and displaying information about interfaces that are operating in channelized mode, you must specify the channel number in the interface name (*physical:channel*); for example, `so-2/2/0:0` and `so-2/2/0:1`. For more information about interface names, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Concatenated and Nonconcatenated Mode

On SONET OC-48 interfaces that are configured for channelized (multiplexed) mode, the bytes e1-quiet and bytes f1 options in the sonet-options statement have no effect. The bytes f2, bytes z3, bytes z4, and path-trace options work correctly on channel 0. These bytes work in the transmit direction only on channels 1, 2, and 3.

The M160 four-port POS OC-12 PIC can run each of the OC-12 links in concatenated mode only and requires a Type 2 M160 Flexible PIC Concentrator (FPC). Similarly, the four-port POS OC-3 PIC cannot run in nonconcatenated mode on any platform.

Channelized DS-3 to DS-0 Naming

You can configure 28 T1 channels per T3 interface. Each T1 link can have up to eight channel groups, and each channel group can hold any combination of DS-0 timeslots. To specify the T1 link and DS-0 channel group number in the name, use colons (:) as separators. For example, a Channelized DS-3 to DS-0 PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x:y
```

where *x* is a T1 link ranging from 0 through 27 and *y* is a DS-0 channel group ranging from 0 through 7 (see Table 21, “Ranges for Channelized DS-3 to DS-0 Configuration” on page 491 for more information about ranges).

You can use any of the values within the range available for *x* and *y*; you do not have to configure the links sequentially. The software applies the interface options you configure according to the following rules:

You can configure t3-options for t1 link 0 and channel group 0 only; for example, ds-0/0/0:0:0.

You can configure t1-options for any t1 link value, but only for channel group 0; for example, ds-0/0/0:x:0.

There are no restrictions on changing the default ds0-options.

If you delete a configuration you previously committed for channel group 0, the options return to the default values.

To configure the channel groups and timeslots for a Channelized DS-3 interface, include the channel-group and timeslots statements at the [edit chassis fpc slot-number pic *pic-number* ct3 port *port-number* t1 *link-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
fpc slot-number {
  pic pic-number {
    ct3 {
      port port-number {
        t1 link-number {
          channel-group group-number timeslots slot-number;
        }
      }
    }
  }
}
```

}



If you commit the interface name but do not include the [edit chassis] configuration, the Channelized DS-3 to DS-0 PIC behaves like a Channelized DS-3 to DS-1 PIC: none of the DS-0 functionality is accessible.

Note

Table 21 shows the ranges for each of the quantities in the preceding configuration.

Table 21: Ranges for Channelized DS-3 to DS-0 Configuration

Item	Variable	Range
FPC slot	<i>slot-number</i>	0 through 7 (see note below)
PIC slot	<i>pic-number</i>	0 through 3
Port	<i>port-number</i>	0 through 1
T1 link	<i>link-number</i>	0 through 27
DS-0 channel group	<i>group-number</i>	0 through 7
timeslot	<i>slot-number</i>	1 through 24



FPC slot range depends on platform. The maximum range of 0 through 7 applies to M40 routers; for M20 routers, the range is 0 through 3; for M10 routers the range is 0 through 1; for M5 routers, the only applicable value is 0. (The Multichannel DS-3 (Channelized DS-3 to DS-0) PIC is not supported on M160 routers.)

Note

Bandwidth limitations restrict the interface to a maximum of 128 channel groups per T3 port, rather than the theoretical maximum of $8 \times 28 = 224$.

There are 24 timeslots on a T1 interface. You can designate any combination of timeslots for usage, but you can use each timeslot number on only one channel group within the same T1 link.

To use timeslots 1 through 10, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number timeslots 1-10;
```

To use timeslots 1 through 5, timeslot 10, and timeslot 24, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number timeslots 1-5,10,24;
```

Note that spaces are not allowed when you specify timeslot numbers. For more information about these interfaces, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Channelized E1 Naming

Each Channelized E1 PIC has 10 E1 ports that you can channelize to the $N \times DS-0$ level. Each E1 interface has 32 timeslots (DS-0), in which timeslot 0 is reserved. You can combine one or more of these DS-0 (channels) to create a channel group ($N \times DS-0$). There can be a maximum of 24 channel groups per E1 interface. Thus, you can configure as many as 240 channel groups per PIC (10 ports \times 24 channel groups per port).

To specify the DS-0 channel group number in the interface name, include a colon (:) as a separator. For example, a Channelized E1 PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x
```

where x is a DS-0 channel group ranging from 0 through 23 (see Table 22 on page 493 for more information about ranges).

You can use any of the values within the range available for x ; you do not have to configure the links sequentially. The software applies the interface options you configure according to the following rules:

You can configure the `e1-options` statement for channel group 0 only; for example, `ds-0/0/0:0`.

There are no restrictions on changing the default `ds0-options`.

If you delete a configuration you previously committed for channel group 0, the options return to the default values.

To configure the channel groups and timeslots for a Channelized E1 interface, include the `channel-group` and `timeslots` statements at the `[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
  fpc slot-number {
    pic pic-number {
      ce1 {
        e1 port-number {
          channel-group group-number timeslots slot-number;
        }
      }
    }
  }
}
```



Note

If you commit the interface name but do not include the `[edit chassis]` configuration, the Channelized E1 PIC behaves like a standard E1 PIC: none of the DS-0 functionality is accessible.

Table 22 shows the ranges for each of the quantities in the preceding configuration.

Table 22: Ranges for Channelized E1 Configuration

Item	Variable	Range
FPC slot	<i>slot-number</i>	0 through 7 (see note below)
PIC slot	<i>pic-number</i>	0 through 3
E1 port	<i>port-number</i>	0 through 9
DS-0 channel group	<i>group-number</i>	0 through 23
timeslot	<i>slot-number</i>	1 through 32

**Note**

FPC slot range depends on platform. The maximum range of 0 through 7 applies to M40 routers; for M20 routers, the range is 0 through 3; for M10 routers the range is 0 through 1; for M5 routers, the only applicable value is 0. (The Channelized E1 PIC is not supported on M160 routers.)

The theoretical maximum number of channel groups possible per PIC is $10 \times 24 = 240$. This is within the maximum bandwidth available.

There are 32 timeslots on an E1 interface. You can designate any combination of timeslots for usage.

To use timeslots 1 through 10, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
channel-group group-number timeslots 1-10;
```

To use timeslots 1 through 5, timeslot 10, and timeslot 24, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
channel-group group-number timeslots 1-5,10,24;
```

Note that spaces are not allowed when you specify timeslot numbers.

For further information about these interfaces, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Channelized STM-1 Interface Virtual Tributary Mapping

By default, virtual tributary mapping uses KLM mode. You can configure virtual tributary mapping to use KLM or ITU-T mode. On the original Channelized STM-1 PIC, to configure virtual tributary mapping, include the *vtmapping* statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
vtmapping (klm | itu-t);
```

For the Channelized STM-1 PIC with QPP, you can configure virtual tributary mapping by including the *vtmapping* statement at the [edit interfaces cau4 fpc slot-number pic pic-number sonet-options] hierarchy level. For more information, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Configure Layer 2 Circuit Cell-Relay or AAL5 Transport Mode

By default, ATM layer 2 circuit mode is set to AAL5, which tunnels IP packets over an ATM backbone. To tunnel a stream of AAL5-encoded ATM cells over an IP MPLS backbone, use Layer 2 circuit AAL5 transport mode. To tunnel a stream of ATM cells over an IP MPLS backbone, use Layer 2 circuit cell relay transport mode.

To configure the Layer 2 circuit transport modes for ATM2 only, include the `atm-l2circuit-mode` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
atm-l2circuit-mode (cell | aal5);
```

`aal5` tunnels a stream of AAL5-encoded ATM cells over an IP backbone

`cell` tunnels a stream of ATM cells over an IP backbone.

For more information about ATM Martini transport mode, see *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service* and *JUNOS Internet Software Feature Guide*.

Configure the Drop Policy for Traffic with Source-Route Constraints

By default, the router forwards IP traffic that has either loose or strict source-route constraints. However, you might want the router to use only the IP destination address on transit traffic for forwarding decisions. You can configure the router to discard IP traffic with source-route constraints by including the `no-source-route` statement at the `[edit chassis]` hierarchy level:

```
[edit chassis]
no-source-route;
```

Configure Redundancy

For routers that have multiple Routing Engines or multiple System and Switch Boards (SSBs), you can configure redundancy properties. A separate log file is provided for redundancy logging, located at `/var/log/mastership`.

This section describes the following tasks for configuring redundancy:

Configure Routing Engine Redundancy on page 495

Default Routing Engine Redundancy Behavior on page 500

Configure SFM Redundancy on page 501

Configure SSB Redundancy on page 501

For information about how to synchronize Routing Engines, see “Synchronize Routing Engines” on page 191.

Configure Routing Engine Redundancy

For routers with two Routing Engines, you can configure which Routing Engine is the master and which is the backup. By default, the Routing Engine in slot 0 is the master (RE0) and the one in slot 1 is the backup (RE1).

To modify the default configuration, include the routing-engine statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
routing-engine slot-number (master | backup | disabled);
```

slot-number can be 0 or 1. To configure the Routing Engine to be the master, specify the master option. To configure it to be the backup, specify the backup option. To switch between the master and the backup Routing Engines, you must modify the configuration and then activate the configuration by issuing the commit command.



Note

We recommend that both Routing Engines have the same configuration.

You can use either the console port or the management Ethernet (fxp0) port to establish connectivity between the two Routing Engines. You can then copy or ftp the configuration from the master to the backup, and load the file and commit it in the normal way.

To make a tty connection to the other Routing Engine using the router's internal Ethernet network, issue the following command:

```
user@host > request routing-engine login (other-routing-engine | re0 | re1)
```

You can configure Routing Engine redundancy in the following ways:

Copy a Configuration File from One Routing Engine to the Other on page 496

Load a Package from the Other Routing Engine on page 497

Change over to the Backup Routing Engine if it Detects a loss of KeepAlive Signal on page 497

Change to the Backup Routing Engine without Interruption to Packet Forwarding (Graceful Switchover) on page 499

Copy a Configuration File from One Routing Engine to the Other

To copy a configuration file from one Routing Engine to the other, you use the existing file copy command:

```
user@host > file copy source destination
```

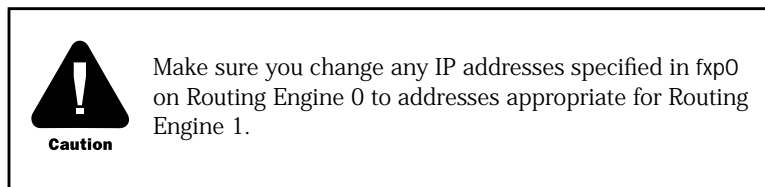
In this case, *source* is the name of the configuration file. These files are stored in the directory /config. The active configuration is /config/juniper.conf, and older configurations are in /config/juniper.conf { 1...9 }. *destination* is a file on the other Routing Engine.

The following is an example of copying a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

To load the file into configuration mode, use the load replace configuration mode command:

```
user@host% load replace /var/tmp/copied-juniper.conf
```



You can use configuration groups to ensure that the correct IP addresses are used for each Routing Engine and to maintain a single configuration file for both Routing Engines.

The following example defines configuration groups re0 and re1 with separate IP addresses. These well-known configuration group names take effect only on the appropriate Routing Engine.

```
groups {
  re0 {
    system {
      host-name my-re0;
    }
    interfaces {
      fxp0 {
        description "10/100 Management interface";
        unit 0 {
          family inet {
            address 10.255.2.40/24;
          }
        }
      }
    }
  }
}
```

```

re1 {
  system {
    host-name my-re1;
  }
  interfaces {
    fxp0 {
      description "10/100 Management interface";
      unit 0 {
        family inet {
          address 10.255.2.41/24;
        }
      }
    }
  }
}

```

For more information about the configuration groups feature, see “Configuration Groups” on page 213.

Load a Package from the Other Routing Engine

You can load a package from the other Routing Engine onto the local Routing Engine using the existing request system software add *package-name* command:

```
user@host > request system software add re(0|1):/filename
```

In the *re* portion of the URL, specify the number of the other Routing Engine. In the *filename* portion of the URL, specify the path to the package. Packages are typically in the directory `/var/sw/pkg`.

Change over to the Backup Routing Engine if it Detects a loss of KeepAlive Signal

Once you have configured a backup Routing Engine, you can direct it to assume mastership automatically if it detects loss of keepalive signal from the master. By default, this feature is disabled; to enable it, include the failover on-loss-of-keepalives statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
failover on-loss-of-keepalives;
```

By default, failover will occur after 300 seconds (5 minutes). To change the keepalive time period, include the keepalive-time statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
keepalive-time seconds;
```

The range for keepalive-time is 2 through 10,000 seconds.

If you configure the keepalive time for 2 seconds, the sequence of events is as follows:

1. After 2 seconds of keepalive loss, a message is logged.
2. After 2 seconds of keepalive loss, the backup Routing Engine attempts to assume mastership. An alarm is generated whenever the backup is active and the display is updated with current status.

3. Once the backup Routing Engine assumes mastership, it continues to function as master even after the originally configured master Routing Engine has successfully resumed operation. You must intervene to restore its previous backup status. However, if at any time one of the Routing Engines is not present, the other one becomes master automatically, regardless of how redundancy is configured.



Note

Packet forwarding is interrupted when you enable this feature. To configure a backup Routing Engine to assume mastership automatically without any interruption to packet forwarding, see “Change to the Backup Routing Engine without Interruption to Packet Forwarding (Graceful Switchover)” on page 499.

Change to the Backup Routing Engine without Interruption to Packet Forwarding (Graceful Switchover)

For routers with two Routing Engines, you can configure graceful switchover. When you enable this feature, the backup Routing Engine automatically synchronizes its configuration and state with the master Routing Engine. Any update to the master Routing Engine state is replicated on the backup Routing Engine. When the backup Routing Engine assumes mastership, the Packet Forwarding Engine (PFE) deletes its connection with the old master Routing Engine and reconnects with the new master Routing Engine. If the new master Routing Engine detects that the PFE state is not up to date, it resends state update messages.

The master Routing Engine sends periodic keepalives to the backup Routing Engine. If the backup Routing Engine does not receive a keepalive after 2 seconds (the default value) from the master Routing Engine, it assumes that the master Routing Engine has failed and assumes mastership without interruption to packet forwarding. The backup Routing Engine does not receive a keepalive signal when the master Routing Engine has failed or is removed. If this happens, the backup Routing Engine assumes mastership. When you reboot the master Routing Engine, mastership switches over to the backup Routing Engine. For more information about keepalive signals, see “Change over to the Backup Routing Engine if it Detects a loss of KeepAlive Signal” on page 497.

When you enable graceful switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files—for example, accounting and traceoptions, are not replicated to the backup Routing Engine. Local statistics—for example, rpd and lsp statistics are not maintained.

When graceful switchover occurs, some offline FRUs might come online. On T-series routing platforms, SIBS restart one at a time.

If you modify the configuration after you have enabled graceful switchover, you must issue the `commit synchronize` command to synchronize both Routing Engines. We recommend issuing the `commit synchronize` command on the master Routing Engine. If you issue this command on the backup Routing Engine, the JUNOS software displays a warning and commits the candidate configuration. You cannot issue the `commit` command without the `commit synchronize` option after you enabled graceful switchover. If you issue this command, the JUNOS software displays a warning. For information about the `commit synchronize` command, see “Synchronize Routing Engines” on page 191.

A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine. When you enable graceful switchover, the CLI indicates which Routing Engine you are using. For example:

```
{master}[edit]
user@host#
```

**Note**

You must use the same version of JUNOS software on both Routing Engines. If you are performing a software upgrade, disable graceful switchover.

To enable switchover when a daemon fails, include the `other-routing-engine` option at the [edit system processes *process-name* failure] hierarchy level. For example, if you want graceful switchover to take place when the routing process fails, include the `other-routing-engine` option at the [edit system processes routing failure] hierarchy level.

You must enable graceful restart on all the protocols you have configured at the [edit protocols] hierarchy level. If you have configured a protocol that does not support graceful restart, graceful switchover might not work.

For information about graceful restart, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

By default, graceful switchover is disabled. To enable it, include the `graceful-switchover` statement and specify `enable` at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
graceful-switchover (disable | enable)
```

Default Routing Engine Redundancy Behavior

By default, the JUNOS software uses `re0` as the master Routing Engine and `re1` as the backup Routing Engine. Unless otherwise specified in the configuration, `re0` will always assume mastership if the acting-master Routing Engine is rebooted.

Take the following steps to see how the default routing engine redundancy setting works:

1. Make sure the router is running on `re0` as the master Routing Engine.
2. Manually switch the state of Routing Engine mastership. `re0` is now the backup Routing Engine and `re1` is the master Routing Engine. For information about switching routing engine mastership, see the *JUNOS Internet Software Configuration Guide: Operational Mode Command Reference*.

**Note**

On the next reboot of the master Routing Engine, the JUNOS software returns the router to the default state because you have not configured the Routing Engines to maintain this state after a reboot.

3. Reboot the master Routing Engine `re1`. When you do this, the Routing Engine boots up and reads the configuration. Because you have not specified in the configuration which Routing Engine is the master, `re1` uses the default configuration as the backup. Now both `re0` and `re1` are in a backup state. The JUNOS software detects this conflict and, to prevent a no-master state, reverts to the default configuration to direct `re0` to assume mastership.

Configure SFM Redundancy

For routers with two Switching and Forwarding Modules (SFM), you can configure which SFM is the master and which is the backup. By default, the SFM in slot 0 is the master and the one in slot 1 is the backup. To modify the default configuration, include the `sfm` statement at the `[edit chassis redundancy]` hierarchy level:

```
[edit chassis redundancy]
  sfm slot-number (always | preferred);
```

`slot-number` can be 0 or 1.

`always` defines the SFM as the sole device.

`preferred` defines a preferred SFM.



Note

SFM redundancy is for M40e and M160 routers only.

Configure SSB Redundancy

For routers with two System and Switch Boards (SSB), you can configure which SSB is the master and which is the backup. By default, the SSB in slot 0 is the master and the one in slot 1 is the backup. To modify the default configuration, include the `ssb` statement at the `[edit chassis redundancy]` hierarchy level:

```
[edit chassis redundancy]
  ssb slot-number (always | preferred);
```

`slot-number` can be 0 or 1.

`always` defines the `ssb` as the sole device.

`preferred` defines a preferred `ssb`.



Note

SSB redundancy is for M20 routers only.

Configure Packet Scheduling

Packet scheduling is for M160 routers only. By default, packet scheduling is disabled. To configure a router to operate in packet-scheduling mode, include the packet-scheduling statement at the [edit chassis] hierarchy level:

```
[edit chassis]
packet-scheduling;
```

To explicitly disable the packet-scheduling statement, include the no-packet-scheduling statement at the [edit chassis] hierarchy level:

```
[edit chassis]
no-packet-scheduling;
```

When you enable packet-scheduling mode, the Packet Director ASIC schedules packet dispatches to compensate for transport delay differences. This preserves the interpacket gaps as the packets are distributed from the Packet Director ASIC to the Packet Forwarding Engine.

Whenever you change the configuration for packet-scheduling, the system stops all Switching and Forwarding Modules (SFMs) and Flexible Pic Concentrators (FPCs) and restarts them in the new mode.

Configure the Link Services PICs

The Multilink Protocol enables you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.

The Link Services PIC supports the following Multilink Protocol encapsulation types at the logical unit level:

Multilink Point-to-Point Protocol (MLPPP)

Multilink Frame Relay (MLFR FRF.15)

The Link Services PIC also supports the Multilink Frame Relay User-to-Network Interface (UNI) and the Network-to-Network Interface (NNI) (MLFR FRF.16) encapsulation type at the physical interface level.

MLFR (FRF.16) is supported on a channelized interface, *ls-fpc/pic/port:channel*, which denotes a single MLFR (FRF.16) bundle. For MLFR (FRF.16), multiple links are combined to form one logical link. Packet fragmentation and reassembly occur on a per-VC basis. Each bundle can support multiple VCs. The physical connections must be E1, T1, Channelized DS-3 to DS-1, Channelized DS-3 to DS-0, Channelized E1, Channelized STM-1, or Channelized QPP interfaces.

The default number of bundles per Link Services PIC is 16, ranging from *ls-fpc/pic/port:0* to *ls-fpc/pic/port:15*.

To configure the number of bundles on a Link Services PIC, include the *mlfr-uni-nni-bundles* statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
mlfr-uni-nni-bundles number;
```

The maximum number of MLFR UNI NNI bundles each Link Services PIC can accommodate is 128. A link can associate with one link services bundle only. For more information, see the *JUNOS Internet Software Guide: Network Interfaces and Class of Service* .



Note

The Link Services PIC is not compatible with the M160 or T-series platforms.

Configure the Idle Cell Format

ATM devices send idle cells to enable the receiving ATM interface to recognize the start of each new cell. The receiving ATM device does not act on the contents of idle cells and does not pass them up to the ATM layer in the ATM protocol stack.

By default, the idle cell format for ATM cells is (4 byte): 0x00000000. For ATM 2 PICs only, you can configure the format of the idle cell header and payload bytes.

To configure the idle cell header to use the International Telecommunications Union (ITU-T) standard of 0x00000001, include the `itu-t` statement at the `[edit chassis fpc number pic number idle-cell-format]` hierarchy level:

```
[edit chassis fpc number pic number idle-cell-format]
itu-t;
```

By default, the payload pattern is cell payload (48 bytes). To configure the idle cell payload pattern, include the `payload-pattern` statement at the `[edit chassis fpc number pic number idle-cell-format]` hierarchy level:

```
[edit chassis fpc number pic number idle-cell-format]
payload-pattern payload-pattern-byte;
```

The payload pattern byte can range from 0x00 through 0xff.

Configure Path MTU Check

By default, MTU check is disabled. On M-series routers, you can configure path MTU checks on the outgoing interface for unicast traffic routed on a VRF. To enable path checks on the outgoing interface for unicast traffic routed on a VRF, include the `mtu-check` statement at the `[edit chassis]` hierarchy level:

To enable MTU check, include the `mtu-check` statement at the `[edit chassis]` hierarchy level:

```
[edit chassis]
vrf-mtu-check;
```

