Juniper Networks | Engineering Simplicity

Junos® OS

Chassis-Level User Guide

Published
2024-12-24

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos® OS Chassis-Level User Guide
Copyright © 2024 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide

Use this guide to configure several properties of a router at the `[edit chassis]` hierarchy level. You can also configure support for chassis-level alarms, power management, and other features at the chassis level. Some of these features are platform-specific, while others are common across all routers.

1

CHAPTER

Overview

Chassis-Level Features Overview

The Junos Software enables you to configure several properties of the router and many PIC-level features at the `[edit chassis]` hierarchy level. Some of the features are specific to M Series, MX Series, J Series, or T Series routers, while some others are common across all routers.

NOTE: Statements at the `[edit chassis redundancy]` hierarchy level are described in the *JUNOS High Availability Configuration Guide*.

NOTE: The configuration statements at the `[edit chassis lcc]` hierarchy level apply only to a routing matrix based on a TX Matrix router or a TX Matrix Plus router. For information about a routing matrix composed of a TX Matrix router and T640 routers, see TX Matrix Router and T640 Router Configuration Overview and the *TX Matrix Router Hardware Guide*. For information about a routing matrix composed of a TX Matrix Plus router and T1600 or T4000 routers, see TX Matrix Plus Router Configuration Overview and the *TX Matrix Plus Router Hardware Guide*.

NOTE: The `sanity-poll` configuration statements at the `[edit chassis fpc slot-number]` hierarchy level apply only to T Series routers. You can also configure `sanity-poll` for routing matrix based on a TX Matrix router or TX Matrix Plus router at the hierarchy level `[edit chassis lcc number fpc number]`.

2

CHAPTER

MPC-specific

Upgrading or Downgrading MPCs (MPC8E and JNP10K-LC2101)

IN THIS SECTION

You can upgrade MPC8E to provide an increased bandwidth of 1600 Gbps (1.6 Tbps). Similarly, you can downgrade the JNP10K-LC2101 MPC to provide a decreased bandwidth of 1.44Tbps.

Upgrading MPC8E to Provide Increased Bandwidth

In Junos OS Release 15.1F5 and later, MX2020 and MX2010 routers support MPC8E (model number: MX2K-MPC8E), a new Modular Port Concentrator (MPC) with two Modular Interface Card (MIC) slots that provide a maximum bandwidth of 960 Gbps. . Each MIC slot on MPC8E supports a 12-port rate selectable MIC (MIC-MRATE). MPC8E has four Packet Forwarding Engines, each providing a maximum bandwidth of 240 Gbps.

In Junos OS Release 16.1R1 and later, you can upgrade MPC8E to provide an increased bandwidth of 1600 Gbps (1.6 Tbps) by using an add-on license. After you purchase the license and perform the upgrade, MPC8E provides a bandwidth of 1.6 Tbps, which is equivalent to the bandwidth that MPC9E provides. However, the MPC continues to be identified as MPC8E.

> **NOTE**: After you upgrade MPC8E to provide a bandwidth of 1.6 Tbps, the power consumption by MPC8E increases and is equivalent to the power that MPC9E consumes. See MPC8E for more information.

After you purchase the add-on license, you upgrade the bandwidth by using the `set chassis fpc slot bandwidth 1.6T` command. You can disable this feature by using the `delete chassis fpc slot bandwidth 1.6T` command.

NOTE: When you modify the bandwidth of MPC8E and commit the configuration, the MPC automatically reboots . For instance, if you upgrade MPC8E to provide a bandwidth of 1.6 Tbps and commit the configuration, the MPC automatically reboots. Also, if you disable the feature on MPC8E (to provide a bandwidth of 960 Gbps) and commit the configuration, the MPC automatically reboots.

Downgrading JNP10K-LC2101 to Provide Decreased Bandwidth

In Junos OS Release 18.2R1 and later, MX10008 routers support the JNP10K-LC2101 line card. The JNP10K-LC2101 is a fixed-configuration MPC that provides increased port density and performance to the MX10008 routers. JNP10K-LC2101 does not contain separate slots for Modular Interface Cards (MICs). MX10008 routers support eight JNP10K-LC2101 MPCs. By default, each JNP10K-LC2101 MPC provides a maximum bandwidth of 2.4Tbps. JNP10K-LC2101 has six Packet Forwarding Engines, each providing a maximum bandwidth of up to 400 Gbps, which cannot be oversubscribed.

You can downgrade the JNP10K-LC2101 MPC to provide a decreased bandwidth of 1.44Tbps. After you perform the downgrade, JNP10K-LC2101 provides a bandwidth of 1.44Tbps. Each of the six Packet Forwarding Engines now provide a maximum bandwidth of up to 240 Gbps, which cannot be oversubscribed.

You can downgrade the bandwidth by using the `set chassis fpc slot bandwidth 1.44T` command. You can disable this feature by using the `delete chassis fpc slot bandwidth 1.44T` command.

NOTE: When you modify the bandwidth of JNP10K-LC2101 and commit the configuration, the MPC automatically reboots . For instance, if you downgrade JNP10K-LC2101 to provide a bandwidth of 1.44 Tbps and commit the configuration, the MPC automatically reboots. Also, if you disable the feature on JNP10K-LC2101 (to provide a bandwidth of 2.4Tbps) and commit the configuration, the MPC automatically reboots.

RELATED DOCUMENTATION

No Link Title

Line card (MX10K-LC2101)

Load based throttling for AFT based linecards on MX10004 and MX10008

SUMMARY

The load-based throttling feature, enabled by default, for AFT based linecard MX10K-LC9600 on MX10004 and MX10008, prevents saturation of line card processing capacity, reduces programming delays, and improves efficiency. The Packet Forwarding Engine (PFE) supports multi-threading and guides the Routing Engine (RE) to control packet management and load balancing. Ping based throttling is supported for non-AFT based line cards. This feature is supported for integrated and disaggregated border network gateway (BNG) modes, on the following interface types:

- Gigabit Ethernet/Line Termination (GE/LT) managed concurrently on single AFT card interface.
- Aggregated Ethernet/Remote Link Termination (AE/RLT) interface facilitates communication and control across various components in a network system on multiple AFT cards.
- AE/RLT interface with AFT and non-AFT cards.

This feature supports the following commands under `[edit] system services resource-monitor` hierarchy:

- The `no-load-throttle` command disables line card load-based throttling.
- In disaggregated BNG mode, the `no-usage-report` statement disables PFCP (Packet Forwarding Control Protocol) usage reports.

See no-load-throttle.

Configuring Line Card Interoperability

IN THIS SECTION

Interoperability of Type 3 FPCs and Type 4 FPCs with Type 5 FPCs

Support for interoperability of Type 3 FPCs, Type 4 FPCs, and T640-FPC4-1P-ES with Type 5 FPCs is now possible with fabric notification translation. This feature is supported on T4000 routers.

Basic packet forwarding, IPv4, IPv6, MPLS, and multicast (dataplane) are currently supported through this feature.

Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode

To configure a 100-Gigabit Ethernet MIC (MIC3-3D-1X100GE-CFP) to interoperate with Juniper Networks Type 4 100-Gigabit Ethernet PICs (model number PD-1CE-CFP-FPC4), you can use the `forwarding-mode` statement with the `sa-multicast` option at the `[edit chassis fpc slot pic slot]` hierarchy level.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the 100-Gigabit Ethernet MIC. The egress packet flow is the traffic flowing from the 100-Gigabit Ethernet MIC to the 100-Gigabit Ethernet PIC. Since no VLAN tags are available, the SA multicast bit is sent on the outgoing packets. At the other end, the 100-Gigabit Ethernet PIC looks at the bit and forwards the packets to either Packet Forwarding Engine 0 or 1. The ingress packet flow is the traffic flowing from a 100-Gigabit Ethernet PIC to a 100-Gigabit Ethernet MIC. When the 100-Gigabit Ethernet PIC is sending out a packet, the multicast bit is set based on the Packet Forwarding Engine packet received. The multicast bit is then transmitted and the MPC3E sees the multicast bit on ingress.

NOTE: The SA multicast bit is ignored by MPC3E while learning the source MAC addresses.

Configuring 100-Gigabit Ethernet MICs

The interoperability mode between the 100-Gigabit Ethernet MIC and the 100-Gigabit Ethernet PIC is configured on a PIC basis. The MPC3E has two MIC slots. A 100-Gigabit Ethernet MIC installed in slot 0 corresponds to `pic 0`, and the MIC installed in slot 1 corresponds to `pic 2`.

NOTE: The configuration is valid only on PIC 0 and PIC 2.

To configure SA multicast mode on a Juniper Networks 100-Gigabit Ethernet MIC in MPC 0, PIC 0 for interconnection with another Juniper Networks 100-Gigabit Ethernet PIC, use the `set chassis fpc slot pic slot forwarding-mode sa-multicast` command, as follows:

```
[edit chassis fpc slot pic slot]
forwarding-mode {
    sa-multicast;
}
```

You can use the `show forwarding-mode` command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
```

Configuring 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4)

The default packet steering mode for the 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4)is SA multicast bit mode. There is no SA multicast configuration required on the 100-Gigabit Ethernet PIC to enable this mode.

NOTE: SA multicast mode can be configured, but it is not necessary.

The 100-Gigabit Ethernet PIC uses a Type 4 FPC and two 50 Gpbs Packet Forwarding Engines to achieve 100 Gbps throughput. The 50 Gpbs physical interfaces are created when the 100-Gigabit Ethernet PIC is installed. The two physical interfaces are visible and configuration is allowed on both physical interfaces. The physical interfaces on the 100-Gigabit Ethernet PIC should be configured in static LAG mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet MIC instead of two independent 50 Gpbs interfaces.

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one AE physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces `et-`*`fpc/pic/0:0`* and `et-`*`fpc/pic/0:1`* where *fpc* is the FPC slot number and *pic* is the PIC slot number. The example shows how to configure two physical interfaces for PIC 0 in FPC 5:

```
chassis {
    aggregated-devices {
        ethernet {
            device-count 1;
        }
    }
}
interfaces {
    et-5/0/0:0 {
        gigether-options {
            802.3ad ae0;
        }
    }
    et-5/0/0:1 {
        gigether-options {
```

```
            802.3ad ae0;
        }
    }
}
```

Interoperability Between MPC4E (MPC4E-3D-2CGE-8XGE) and 100-Gigabit Ethernet PICs on Type 4 FPC

You can enable interoperability between the MPC4E (MPC4E-3D-2CGE-8XGE) and the 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4) by:

- Enabling source address (SA) multicast bit steering mode on the MPC4E.
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the MPC4E. The egress packet flow is the traffic flowing from the MPC4E to the 100-Gigabit Ethernet PIC. Because no VLAN tags are available, the SA multicast bit is sent on the outgoing packets. At the other end, the 100-Gigabit Ethernet PIC checks the multicast bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. The ingress packet flow is the traffic flowing from the 100-Gigabit Ethernet PIC to the MPC4E. When the 100-Gigabit Ethernet PIC sends out a packet, the multicast bit is set based on the packet received from the Packet Forwarding Engine. The multicast bit is then transmitted and the MPC4E checks the multicast bit on ingress.

The 100-Gigabit Ethernet PIC uses a Type 4 FPC and two 50-Gbps Packet Forwarding Engines to achieve a throughput of 100 Gbps. The 50-Gbps physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both physical interfaces. The physical interfaces on the 100-Gigabit Ethernet PIC must be configured in static LAG mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit Ethernet aggregated interface is visible on the link connecting to the MPC4E instead of two independent 50-Gbps interfaces.

Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

IN THIS SECTION

You can enable interoperability between the MPC4E and the 100-Gigabit Ethernet PIC by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on MPC4E

The interoperability mode between the MPC4E and the 100-Gigabit Ethernet PIC is configured on a PIC basis. MPC4E-3D-2CGE-8XGE is a fixed-configuration MPC and does not contain separate slots for Modular Interfaces Cards (MICs). MPC4E contains two Packet Forwarding Engines—**PFE 0** hosts **PIC 0** and **PIC 1** and **PFE 1** hosts **PIC 2** and **PIC 3**.

NOTE: This configuration is valid only on **PIC 1** and **PIC 3**.

To configure SA multicast mode on **PIC 1** of an MX480 router with MPC4E for interconnection with the 100-Gigabit Ethernet PIC:

1. To specify the forwarding mode as `sa-multicast`, include the `forwarding-mode` statement at the `[edit chassis fpc` *`slot`* `pic` *`slot`*`]` hierarchy level.

```
[edit chassis]
user@host # set fpc 3 pic 1 forwading-mode sa-multicast
```

2. To verify that the forwarding mode is set to `sa-multicast`, issue the following command:

```
[edit chassis fpc 3 pic 1]
user@host # show forwarding-mode
```

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the Ethernet PIC as One Aggregated Ethernet Interface

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-x/y/0:0 and et-x/y/0:,1 where *x* is the FPC slot number and *y* is the PIC slot number.

The default packet steering mode for the 100-Gigabit Ethernet PIC is SA multicast bit mode. No SA multicast configuration is required on the 100-Gigabit Ethernet PIC to enable this mode.

> **NOTE**: SA multicast mode can be configured, but it is not necessary.

1. To specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis aggregated-devices ethernet]
user@host # set device-count 2
```

2. To specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host # set et-4/3/0:0 gigether-options 802.3ad ae0
user@host # set et-4/3/0:1 gigether-options 802.3ad ae0
```

3. Verify the configuration at the interface.

```
[edit]
user@host # show interfaces
```

```
..
et-4/3/0:0 {
     gigether-options {
  802.3ad ae0;
 }
 }
et-4/3/0:1 {
   gigether-options {
    802.3ad ae0;
```

```
}
}
```

Interoperability Between MPC7E-MRATE and 100-Gigabit Ethernet PICs on Type 4 FPC

You can enable interoperability between the MPC7E (MPC7E-MRATE) and the 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4) by:

- Enabling source address (SA) multicast bit steering mode on the MPC7E
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the MPC7E. The egress packet flow is the traffic flowing from the MPC to the 100-Gigabit Ethernet interface. Because no VLAN tags are available, the SA multicast bit is sent on the outgoing packets. At the other end, the 100-Gigabit Ethernet interface checks the multicast bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. The ingress packet flow is the traffic flowing from the 100-Gigabit Ethernet interface to the MPC7E. When the 100-Gigabit Ethernet interface sends out a packet, the multicast bit is set based on the packet received from the Packet Forwarding Engine. The multicast bit is then transmitted and the MPC7E checks the multicast bit on ingress.

The 100-Gigabit Ethernet PIC uses a Type 4 FPC and two 50-Gbps Packet Forwarding Engines to achieve a throughput of 100 Gbps. The 50-Gbps physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both physical interfaces. The physical interfaces on the 100-Gigabit Ethernet PIC must be configured in static LAG mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit Ethernet aggregated interface is visible on the link connecting to the MPC7E instead of two independent 50-Gbps interfaces.

Configuring MPC7E-MRATE to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

IN THIS SECTION

You can enable interoperability between the MPC7E (MPC7E-MRATE) and the 100-Gigabit Ethernet PIC by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on MPC7E

The interoperability mode between the MPC7E (MPC7E-MRATE) and the 100-Gigabit Ethernet PIC is configured on a PIC basis. MPC7E is a fixed-configuration MPC and does not contain separate slots for Modular Interfaces Cards (MICs). MPC7E contains two Packet Forwarding Engines—**PFE 0** hosts **PIC 0** and **PFE 1** hosts **PIC 1** .

To configure SA multicast mode on `FPC13`,**PIC 1** of MPC7E-MRATE for interconnection with the 100-Gigabit Ethernet PIC:

1. To specify the forwarding mode as `sa-multicast`, include the `forwarding-mode` statement at the `[edit chassis fpc` *`slot`* `pic` *`slot`*`]` hierarchy level.

```
[edit chassis]
user@host # set fpc 13 pic 1 forwading-mode sa-multicast
```

2. To verify that the forwarding mode is set to `sa-multicast`, issue the following command:

```
[edit chassis fpc 13 pic 1]
user@host # show forwarding-mode
```

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the Ethernet PIC as One Aggregated Ethernet Interface

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-x/y/0:0 and et-x/y/0:,1 where *x* is the FPC slot number and *y* is the PIC slot number.

The default packet steering mode for the 100-Gigabit Ethernet PIC is SA multicast bit mode. No SA multicast configuration is required on the 100-Gigabit Ethernet PIC to enable this mode.

NOTE: SA multicast mode can be configured, but it is not necessary.

1. To specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis aggregated-devices ethernet]
user@host # set device-count 2
```

2. To specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host # set et-4/3/0:0 gigether-options 802.3ad ae0
user@host # set et-4/3/0:1 gigether-options 802.3ad ae0
```

3. Verify the configuration at the interface.

```
[edit]
user@host # show interfaces
```

```
..
et-4/3/0:0 {
     gigether-options {
  802.3ad ae0;
 }
 }
et-4/3/0:1 {
   gigether-options {
    802.3ad ae0;
 }
 }
```

Interoperability Between MPC8E (MX2K-MPC8E) and 100-Gigabit Ethernet PICs on Type 4 FPC

You can enable interoperability between the MPC8E (MX2K-MPC8E) and the 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4) by:

- Enabling source address (SA) multicast bit steering mode on the MPC8E.
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the MPC8E. The egress packet flow is the traffic flowing from the MPC to the 100-Gigabit Ethernet Interface. Because no VLAN tags are available, the SA multicast bit is sent on the outgoing packets. At the other end, the 100-Gigabit Ethernet Interface checks the multicast bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. The ingress packet flow is the traffic flowing from the 100-Gigabit Ethernet Interface to the MPC8E. When the 100-Gigabit Ethernet Interface sends out a packet, the multicast bit is set based on the packet received from the Packet Forwarding Engine. The multicast bit is then transmitted and the MPC8E checks the multicast bit on ingress.

The 100-Gigabit Ethernet PIC uses a Type 4 FPC and two 50-Gbps Packet Forwarding Engines to achieve a throughput of 100 Gbps. The 50-Gbps physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both physical interfaces. The physical interfaces on the 100-Gigabit Ethernet PIC must be configured in static LAG mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit Ethernet aggregated interface is visible on the link connecting to the MPC8E instead of two independent 50-Gbps interfaces.

Configuring MPC8E to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

IN THIS SECTION

You can enable interoperability between the MPC8E (MX2K-MPC8E) and the 100-Gigabit Ethernet PIC by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on MPC8E

The interoperability mode between the MPC8E and the 100-Gigabit Ethernet PIC is configured on a PIC basis. MPC8E (MX2K-MPC8E) is a modular MPC that contains two slots for Modular Interfaces Cards (MICs). MPC8E contains four Packet Forwarding Engines—**PIC 0** hosts **PFE 0** and **PFE 1**. **PIC 1**hosts **PFE 2** and **PFE 3**.

To configure SA multicast mode on `FPC 7`,**PIC 1** of MPC8E for interconnection with the 100-Gigabit Ethernet PIC:

1. To specify the forwarding mode as `sa-multicast`, include the `forwarding-mode` statement at the `[edit chassis fpc` *`slot`* `pic` *`slot`*`]` hierarchy level.

```
[edit chassis]
user@host # set fpc 7 pic 1 forwading-mode sa-multicast
```

2. To verify that the forwarding mode is set to `sa-multicast`, issue the following command:

```
[edit chassis]
user@host # show fpc 7 pic 1 forwarding-mode
```

```
sa-multicast;
```

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the Ethernet PIC as One Aggregated Ethernet Interface

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-x/y/0:0 and et-x/y/0:,1 where *x* is the FPC slot number and *y* is the PIC slot number.

The default packet steering mode for the 100-Gigabit Ethernet PIC is SA multicast bit mode. No SA multicast configuration is required on the 100-Gigabit Ethernet PIC to enable this mode.

> **NOTE**: SA multicast mode can be configured, but it is not necessary.

1. To specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis aggregated-devices ethernet]
user@host # set device-count 2
```

2. To specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host # set et-4/3/0:0 gigether-options 802.3ad ae0
user@host # set et-4/3/0:1 gigether-options 802.3ad ae0
```

3. Verify the configuration at the interface.

```
[edit]
user@host # show interfaces
```

```
..
et-4/3/0:0 {
     gigether-options {
  802.3ad ae0;
 }
 }
et-4/3/0:1 {
   gigether-options {
    802.3ad ae0;
 }
 }
```

Interoperability Between MPC9E (MX2K-MPC9E) and 100-Gigabit Ethernet PICs on Type 4 FPC

You can enable interoperability between the MPC9E (MX2K-MPC9E) and the 100-Gigabit Ethernet PIC (PD-1CE-CFP-FPC4) by:

- Enabling source address (SA) multicast bit steering mode on the MPC9E.

- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the MPC9E. The egress packet flow is the traffic flowing from the MPC9E to the 100-Gigabit Ethernet Interface. Because no VLAN tags are available, the SA multicast bit is sent on the outgoing packets. At the other end, the 100-Gigabit Ethernet Interface checks the multicast bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. The ingress packet flow is the traffic flowing from the 100-Gigabit Ethernet Interface to the MPC9E. When the 100-Gigabit Ethernet interface sends out a packet, the multicast bit is set based on the packet received from the Packet Forwarding Engine. The multicast bit is then transmitted and the MPC9E checks the multicast bit on ingress.

The 100-Gigabit Ethernet PIC uses a Type 4 FPC and two 50-Gbps Packet Forwarding Engines to achieve a throughput of 100 Gbps. The 50-Gbps physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both physical interfaces. The physical interfaces on the 100-Gigabit Ethernet PIC must be configured in static LAG mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit Ethernet aggregated interface is visible on the link connecting to the MPC9E instead of two independent 50-Gbps interfaces.

Configuring MPC9E to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

IN THIS SECTION

You can enable interoperability between the MPC9E (MX2K-MPC9E) and the 100-Gigabit Ethernet PIC by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on MPC9E

The interoperability mode between the MPC9E and the 100-Gigabit Ethernet PIC is configured on a PIC basis. MPC9E (MX2K-MPC9E) is a modular MPC that contains two slots for Modular Interfaces Cards

(MICs). MPC9E contains four Packet Forwarding Engines—**PIC 0** hosts **PFE 0** and **PFE 1**. **PIC 1** hosts **PFE 2** and **PFE 3**.

To configure SA multicast mode on `FPC 19`, **PIC 1** of MPC9E for interconnection with the 100-Gigabit Ethernet PIC:

1. To specify the forwarding mode as `sa-multicast`, include the `forwarding-mode` statement at the `[edit chassis fpc slot pic slot]` hierarchy level.

```
[edit chassis]
user@host # set fpc 19 pic 1 forwading-mode sa-multicast
```

2. To verify that the forwarding mode is set to `sa-multicast`, issue the following command:

```
[edit chassis]
user@host # show fpc 19 pic 1 forwarding-mode
```

```
sa-multicast;
```

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the Ethernet PIC as One Aggregated Ethernet Interface

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-x/y/0:0 and et-x/y/0:,1 where *x* is the FPC slot number and *y* is the PIC slot number.

The default packet steering mode for the 100-Gigabit Ethernet PIC is SA multicast bit mode. No SA multicast configuration is required on the 100-Gigabit Ethernet PIC to enable this mode.

NOTE: SA multicast mode can be configured, but it is not necessary.

1. To specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis aggregated-devices ethernet]
user@host # set device-count 2
```

2. To specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host # set et-4/3/0:0 gigether-options 802.3ad ae0
user@host # set et-4/3/0:1 gigether-options 802.3ad ae0
```

3. Verify the configuration at the interface.

```
[edit]
user@host # show interfaces
```

```
..
et-4/3/0:0 {
     gigether-options {
  802.3ad ae0;
 }
 }
et-4/3/0:1 {
   gigether-options {
    802.3ad ae0;
 }
 }
```

RELATED DOCUMENTATION

T4000 FPCs Supported

T4000 PICs Supported

Configuring the Number of Active Ports on 16x10GE 3D MPC

You can disable a sub-set of the physical ports available on the Packet Forwarding Engines of the 16x10GE 3D MPC, and for PICs installed in MPC3, MPC4, MPC5, and MPC6.

Two of the most common reasons for disabling ports are explained below.

- **Ensure guaranteed bandwidth by preventing oversubscription**—The 16x10GE 3D MPC supports one 10-Gigabit Ethernet tunnel interface for each Packet Forwarding Engine. The effective line-rate bandwidth of the MPC is 12 ports because of an oversubscription ratio of 4:3. Therefore, configuring a tunnel interface might further result in the Packet Forwarding Engines being oversubscribed. To prevent such oversubscription and to ensure a guaranteed bandwidth, include the `number-of-ports` configuration statement to disable one or two ports per Packet Forwarding Engine.

- **Enable Switch Control Board (SCB) redundancy**—For maximum bandwidth capabilities (12-port line-rate bandwidth), the 16x10GE 3D MPC uses all the available SCBs (three SCBs for an MX960 router, two SCBs for an MX480 or MX240 router) actively in the chassis.

  If SCB redundancy (2+1 SCBs on an MX960 router or 1+1 SCB on an MX480 or MX240 router) is required, ports on the line card can be disabled by setting the number of usable ports per line card to *8*. In such a case, the third and fourth ports (ports 0/2-3, 1/2-3, 2/2-3, 3/2-3) on every Packet Forwarding Engine are disabled.

To configure the number of active ports on the 16x10GE 3D MPC, include the `number-of-ports` *active-ports* configuration statement at the `[edit chassis fpc slot-number]` hierarchy level:

```
[edit chassis fpc slot-number]
number-of-ports (8 | 12);
```

To configure the number of active ports on a PIC in an MPC3, MPC4, MPC5, or MPC6, include the `number-of-ports` *active-ports* configuration statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
number-of-ports (8 | 12);
```

Specify either 8 or 12 ports using this statement. When eight active ports are configured, two ports per Packet Forwarding Engine are disabled, and the LEDs on the MPC are set to **yellow**. When you specify 12 active ports, one port per Packet Forwarding Engine is disabled and the corresponding LED is set to **yellow**. When you do not include this statement in the configuration, all 16 default ports on the MPC are active.

NOTE:

- Committing the configuration after including the `number-of-ports` *active-ports* configuration statement brings down the Ethernet interfaces for all the ports on the MPC before the ports configuration becomes active.

- A minimum of one high-capacity fan tray is necessary for meeting the cooling requirements of the MPC. The Junos OS generates a chassis **yellow** alarm recommending fan tray upgrade for optimal performance, if the MX router chassis contains an old fan tray.

For more information about the 16x10GE 3D MPC, see the *MX Series Interface Module Reference*.

RELATED DOCUMENTATION

MPC-3D-16XGE-SFPP

Guidelines for Identifying Active PICs on MPC5E-40G10G

MPC5E contains two Packet Forwarding Engines (PFEs) and 4 fixed port PICs. On MPC5E-100G10G, the `PFE0` hosts `PIC0` and `PIC1` while `PFE1` hosts `PIC2` and `PIC3`. All the PICs can be powered on and used.

On the MPC5E-40G10G, the `PFE0` hosts `PIC0` and `PIC2` while `PFE1` hosts `PIC1` and `PIC3`. Only a maximum of two PICs (`PIC0` or `PIC2` and `PIC1` or `PIC3`) can be powered on. The remaining PICs are required to be kept powered off.

This topic describes the guidelines to consider while identifying active PICs on the MPC5E (MPC5E-40G10G):

- By default, (i.e. without any CLI configuration), `PIC0 (12x10GE)` and `PIC1 (12x10GE)` are powered ON while `PIC2 (3x40GE)` and `PIC3 (3x40GE)` shall be powered OFF.
- At least one PIC on every PFE should be configured in power OFF state. `PIC0` and `PIC2` belong to `PFE0` and `PIC1` and `PIC3` belong to `PFE1` .
- If you configure an invalid PIC combination, the default PICs (`PIC0` and `PIC1`) will be powered ON. Also, a syslog message is displayed to indicate the invalid PIC combination selected. When you configure an invalid PIC combination, and commit the change, the commit succeeds and a commit failure message is not displayed.

Table 1 on page 24 lists the active PICs on MPC5E-40G10G based on the configuration.

Table 1: MPC5E-40G10G Active PICs

CLI Configuration	PIC Selection
Default (i.e no CLI configuration)	Online: PIC0 and PIC1 Offline: PIC2 and PIC3
PIC1 ,PIC2 , and PIC3 powered off	Online: PIC0 Offline: PIC1, PIC2, and PIC3
PIC0 ,PIC2 , and PIC3 powered off	Online: PIC1 Offline: PIC0, PIC2, and PIC3
PIC0 ,PIC1 , and PIC3 powered off	Online: PIC2 Offline: PIC0, PIC1, and PIC3
PIC0 ,PIC1 , and PIC2 powered off	Online: PIC3 Offline: PIC0, PIC1, and PIC2
PIC2 and PIC3 powered off	Online: PIC0 and PIC1 Offline: PIC2 and PIC3
PIC1 and PIC2 powered off	Online: PIC0 and PIC3 Offline: PIC1 and PIC2
PIC0 and PIC3 powered off	Online: PIC1 and PIC2 Offline: PIC0 and PIC3
PIC0 and PIC1 powered off	Online: PIC2 and PIC3 Offline: PIC0 and PIC1

Table 1: MPC5E-40G10G Active PICs *(Continued)*

CLI Configuration	PIC Selection
Invalid PIC Configuration (All other combinations of PICs powered off)	Online: `PIC0` and `PIC1` Offline: `PIC2` and `PIC3` **NOTE**: Default PIC configuration is selected for all invalid PIC configurations.

RELATED DOCUMENTATION

6x40GE + 24x10GE MPC5E

3

CHAPTER

Fabric Management

Fabric Resiliency

IN THIS SECTION

Fabric Resiliency and Degradation

IN THIS SECTION

Juniper routers and switches have built in resiliency to tackle failures and error conditions encountered during normal operation. Immediate action is taken by JUNOS software to remedy the failure conditions to minimize traffic loss. No manual intervention is needed. Fabric degradation could be one of the reasons leading to such error conditions. The following sections explain how the PFEs recover in a resilient manner from these failures.

Packet Forwarding Engine Errors and Recovery on PTX Series Routers

Packet Forwarding Engine destinations can become unreachable on PTX Series routers for the following reasons:

- The fabric Switch Interface Boards (SIBs) are offline as a result of a CLI command .
- The fabric SIBs are turned offline by the control board because of high temperature conditions.
- Voltage or polled I/O errors in the SIBs are detected by the control board.
- Unexpected link-training errors occur on all connected planes.
- Two Packet Forwarding Engines can reach the fabric but not each other.
- Link errors occur where two Packet Forwarding Engines have connectivity with the fabric but not through a common plane.

Starting with Junos OS Release 13.3, you can use PTX Series routers to configure Packet Forwarding Engine (PFE)-related error levels and the actions to perform when a specified threshold is reached.

If error levels are not defined, a PTX Series router begins the following phases in the recovery process:

1. SIB restart phase: The router attempts to resolve the issue by restarting the SIBs one by one. This phase does not start if the SIBs are functioning properly and a single line card is facing an issue.
2. SIB and line card restart phase: The router restarts both the SIBs and the line card. If there are line cards that are unable to initiate high-speed links to the fabric after reboot, it is not relevant to loss of live traffic as no interfaces are created for these line cards, preventing the system from issues.
3. Line Card offline phase: Because previous attempts at recovery failed, line cards and interfaces are turned off and the system avoids issues and error conditions.

Fabric Resiliency and Automatic Recovery of Degraded Fabric

Starting Junos Evolved Release 23.4R1, the fabric automatic recovery feature is available to limit data loss. Recovery actions taken include FRU restart, link restart and so on.

The following three-phase fabric recovery actions are attempted at FRU level:

1. FRU level recovery using SIB restart.

2. FRU level recovery using FPC restart or PFE restart.

3. Action for unrecoverable PFEs IFD disable or PFE offline.

> **NOTE**: For platforms that do not have PFE-restart support, FPC restart is provided as the default action.

Fabric recovery action for SIB fault conditions: For reachability faults due to an absent SIB (user driven offline or SIB not present during system power up), Fabric resiliency does not attempt recovery. In systems that do not support fabric recovery, chassis alarms are generated for reachability faults.

PFE Level Recovery Action on PTX Series Routers (PTX10004, PTX10008, and PTX10016 Routers)

For platforms that can support PFE restart, PFE restart will be added as the default phase 2 recovery action.

> **NOTE**: In ASICs with multiple PFEs, the restart affects PPFEs (Per-plane PFEs), similar to PFE offline action.

Recovery decision for phase 2 action is made for either of the following scenarios:

- PFE's with reachability faults all reside in a single FPC.
- PFEs with reachability faults (in one or more FPCs) and have no common of failure.

Phase 2 recovery is attempted on PPFEs that have not recovered from reachability faults after phase 1 recovery.

If the number of PFEs having self reachability faults in an FPC equal to or exceed 50% of the PFEs then the FPC will be restart.

Use the following CLI option to manually configure the default PFE restart action:

```
user@root> set chassis fabric event reachability-fault actions pfe-restart-disable
```

The following table shows the actions on phase 2 recovery, based on the configuration and number of PFEs in fault in an FPC.

Recovery decision	Number of implicated PFEs in FPC	PFE restart supported	PFE restart disable	FPC restart disable	Action
Phase 2 action	<= 50%	Yes	No	x	PFE restart
Phase 2 action	<= 50%	Yes	Yes	No	FPC restart

Phase 2 action	<= 50%	Yes	Yes	Yes	PFE restart
Phase 2 action	>50%	Yes	x	No	FPC restart
Phase 2 action	>50%	Yes	Yes	Yes	PFE restart
Phase 2 action	>50%	Yes	No	Yes	PFE restart

Packet Forwarding Engine Errors and Recovery on T640, T1600 or TX Matrix Routers

Packet Forwarding Engine destinations can become unreachable on T640, T1600 or TX Matrix routers for the following reasons:

- The fabric Switch Interface Boards (SIBs) are offline as a result of a CLI command or a pressed physical button.
- The fabric SIBs are turned offline by the Switch Processor Mezzanine Board (SPMB) because of high temperature conditions.
- Voltage or polled I/O errors in the SIBs are detected by the SPMB.
- All Packet Forwarding Engines receive destination errors on all planes from remote Packet Forwarding Engines, even when the SIBs are online.
- Complete fabric loss is caused by destination timeouts, even when the SIBs are online.

The recovery process consists of the following phases:

1. The router restarts the fabric planes one by one. This phase does not start if the fabric plane is functioning properly and a single line card has issues.
2. Fabric plane and Line Card restart phase: The router restarts both the SIBs and the line cards. If there are line cards that are unable to initiate high-speed links to the fabric after reboot, it is not relevant to loss of live traffic as no interfaces are created for these line cards, preventing the system from issues.
3. Line card offline phase: Because previous attempts at recovery failed, line cards and interfaces are turned off and the system avoids issues and error conditions leading to serious consequences.

NOTE: Starting in Junos OS Release 14.2R6, if a SIB becomes offline because of extreme conditions such as high voltage or high temperature, then as part of the recovery process, the router does not restart the fabric plane for that SIB.

The phased recovery mechanism mentioned above is exhaustive unless there are other errors which could be correlated to these issues.

Starting in Junos OS Release 14.2R6, you can manage fabric degradation in single-chassis systems better by incorporating fabric self-ping and Packet Forwarding Engine liveness mechanisms. Fabric self-ping is a mechanism to detect issues in the fabric data path. Using the fabric self-ping mechanism, every Packet Forwarding Engine ascertains that a packet destined to itself is reaching it when the packet is sent over the fabric path. Packet Forwarding Engine liveness is a mechanism to detect whether a Packet Forwarding Engine is reachable on the fabric plane. To verify that it is reachable, the Packet Forwarding Engine sends a self-destined packet over the fabric plane periodically. If any error is detected by these two mechanisms, the fabric manager raises a *fabric degraded alarm* and initiates recovery by restarting the line card.

MX Series Routers Fabric Resiliency

IN THIS SECTION

MX routers provide intelligent mechanisms to reduce packet loss in hardware failures scenarios. MX Series routers ensure network and service availability with a broad set of multilayered physical, logical, and protocol-level resiliency aspects

MX10008 provides redundancy and resiliency. All major hardware components including the power system, the cooling system, and the control board are fully redundant.

The MX10004 power system and the Routing Control Board (RCB) provide redundancy and resiliency.

The MX2020 and MX2010 chassis provide redundancy and resiliency. All major hardware components including the power system, the cooling system, the control board and the switch fabrics are fully redundant.

Switch Fabric Boards (SFBs) are the data plane for the subsystems in the MX router chassis. SFBs create a highly scalable and resilient “all-active” centralized switch fabric that delivers up to 4 Tbps of full duplex switching capacity to each MPC slot in an MX2000 router.

The MX240, MX480 and MX960 chassis provide redundancy and resiliency. The hardware system is fully redundant, power supplies, fan trays, Routing Engines, and Switch Control Boards.

The MX304 router contains redundant, pluggable, Routing Engines and supports up to three line-card MICs (LMICs).

This topic contains the following sections that describe fabric resiliency options, failure detection methods used, and corrective actions:

Fabric Connectivity Restoration

Packet Forwarding Engine destinations can become unreachable for the following reasons:

- The control boards go offline as a result of a CLI command or a pressed physical button.
- The fabric control boards are turned offline because of high temperature.
- Voltage or polled I/O errors in the fabric.
- All Packet Forwarding Engines receive destination errors on all planes from remote Packet Forwarding Engines, even when the fabrics are online.
- Complete fabric loss caused by destination timeouts, even when the fabrics are online.

When the system detects any unreachable Packet Forwarding Engine destinations, fabric connectivity restoration is attempted. If restoration fails, the system turns off the interfaces to trigger local protection action or traffic re-route on the adjacent routers.

The recovery process consists of the following phases:

1. Fabric plane restart phase: Restoration is attempted by restarting the fabric planes one by one. This phase does not start if the fabric plane is functioning properly and an error is reported by one line card only. An error message is generated to specify that a connectivity loss is the reason for the fabric plane being turned offline. This phase is performed for fabric plane errors only.
2. Fabric plane and line card restart phase: The system waits for the first phase to be completed before examining the system state again. If the connectivity is not restored after the first phase is performed or if the problem occurs again within a duration of 10 minutes, connectivity restoration is attempted by restarting both the fabric planes and the line cards. If you configure the `action-fpc-restart-disable` statement at the `[edit chassis fabric degraded]` hierarchy level to disable restart of the line cards when

a recovery is attempted, an alarm is triggered to indicate that connectivity loss has occurred. In this second phase, three steps are taken:

a. All the line cards that have destination errors on a PFE are turned offline.

b. The fabric planes are turned offline and brought back online, one by one, starting with the spare plane.

c. The line cards that were turned offline are brought back online.

3. Line card offline phase: The system waits for the second phase to be completed before examining the system state again. Connectivity loss is limited by turning the line cards offline and by turning off interfaces because previous attempts at recovery have failed. If the problem is not resolved by restarting the line cards or if the problem recurs within 10 minutes after restarting the line cards, this phase is performed.

The three phases are controlled by timers. During these phases, if an event (such as offlining/onlining line cards or fabric planes) times out, then the phase skips that event and proceeds to the next event. The timer control has a timeout value of 10 minutes. If the first fabric error occurs in a system with two or more line cards, the fabric planes are restarted. If another fabric error occurs within the next 10 minutes, the fabric planes and line cards are restarted. However, if the second fabric error occurs outside of the timeout period of 10 minutes, then the first phase is performed, which is the restart of only the fabric planes.

In cases where all the destination timeouts are traced to a certain line card, for example, one source line card or one destination line card, only that line card is turned offline and online. The fabric planes are not turned offline and online. If another fabric fault occurs within the period of 10 minutes, the line card is turned offline.

By default, the system limits connectivity loss time by detecting severely degraded fabric. No user interaction is necessary.

Line Cards with Degraded Fabric

You can configure a line card with degraded fabric to be moved to the offline state. On an MX10008, MX10004, MX2020, MX2010, MX960, MX480, MX304, or MX240 router, you can configure link errors or bad fabric planes. This configuration is particularly useful in partial connectivity loss scenarios where bringing the line card offline results in faster re-routing. To configure this option on a line card, use the `offline-on-fabric-bandwidth-reduction` statement at the `[edit chassis fpc` *`slot-number`*`]` hierarchy level. For details, see No Link Title, No Link Title, No Link Title, No Link Title, No Link Title and Fabric Plane Management on AS MLC Modular Carrier Card.

Connectivity Loss Towards a Single Destination Only

In certain deployments, a line card indicates a complete connectivity loss towards a single destination only, but it functions properly for other destinations. Such cases are identified and the affected line card is recovered. Consider a sample scenario in which the active planes are 0,1,2,3 and the spare planes are 4,5,6,7 in the connection between line card 0 and line card 1. If line card 0 has single link failures for planes 0 and 1 and if line card 1 has single link failures for planes 2 and 3, a complete connectivity loss occurs between the two line cards. Both line card 0 and line card 1 undergo a phased mode of recovery and fabric healing takes place.

Redundancy Fabric Mode on Active Control Boards

You can configure the active control board to be in redundancy mode or in increased fabric bandwidth mode. To configure redundancy mode for the active control board, use the `redundancy-mode redundant` statement at the `[edit chassis fabric]` hierarchy level.

Detection and Recovery of Fabric-Related Failures Caused by Loss of Connectivity on MX Series Routers

IN THIS SECTION

Connectivity loss in a router occurs when the router is unable to transmit data packets to other neighboring routers, although the interfaces on that router continue to be in the active state. As a result, the other neighboring routers continue to forward traffic to the impacted router, which drops the arriving packets without sending a notification to the other routers.

When a Packet Forwarding Engine in a router is unable to send traffic to other Packet Forwarding Engines over the data plane within the same router, the router is unable to transmit any packets to a neighboring router, although the interfaces are advertised as active on the control plane. Fabric failure can be one of the reasons for the loss of connectivity.

The following fabric failure scenarios can occur:

- Removal of the control board

- High-speed link 2 (HSL2) training failures
- Single link failure on a line card
- Multiple link failures on the same line card or the same fabric plane
- Multiple link failures randomly on a line card or a fabric plane
- Intermittent cyclic redundancy check (CRC) errors
- A complete loss of connectivity for only one destination and not to other destinations

When a line card does not forward traffic due to a certain reason to other line cards within the device, the control protocol on the Routing Engine is unable to detect this condition. The traffic transmission is not diverted to the functional, active line cards and, instead, the packets are continued to be sent to the affected line card and are dropped at that point. The following might be the causes for a line card being unable to forward traffic:

- All the planes in the system are in the `Offline` or `Fault` state.
- All the Packet Forwarding Engines on the line card might have disabled the fabric streams due to destination errors.

If all the Switch Control Boards (SCBs) lose connectivity to the line cards, then all the interfaces are brought down. If a Packet Forwarding Engine of a line card loses complete connectivity to or from the fabric, then that line card is brought down.

System hardware failures can be of the following types:

- A single occurrence or a rare failure for a brief period (such as environmental spikes). This failure is effectively healed without manual intervention by restarting the fabric plane and restarting the line cards and the fabric plane, if necessary.
- Repeated failures that occur frequently.
- A permanent failure.

A recovery from any case of reduced throughput, such as multiple Packet Forwarding Engine destination timeouts on multiple planes is not attempted. Restoration of connectivity is attempted only when all the planes are in the `Offline` or `Fault` state or when the destinations are unreachable on all active planes.

If connectivity loss occurs because of a certain line card, which is either a common source or common destination of the destination timeout, and if you have configured the `action-fpc-restart-disable` statement at the `[edit chassis fabric degraded]` hierarchy level, no recovery action is taken. The `show chassis fabric reachability` command output can be used to verify the status of the fabric and the line card. An alarm is triggered to indicate that the particular line card is causing the connectivity loss.

Fabric-Failure Detection Methods on MX Series Routers

The chassis daemon (chassisd) process detects the removal of a control board. The removal of the control board causes all the active planes that reside on that board to be disabled and a switchover is performed. If the active Routing Engine is also unplugged along with the control board, the detection of the control board removal is delayed until the switchover of the Routing Engine occurs and the reconnection in the primary, backup Routing Engine pair occurs. If the control board is turned offline by specifying the `request chassis cb slot` *`slot-number`* `offline` or a pressed physical button to cause a graceful shutdown, a fabric failure does not occur, even if the control board is moved to the offline state.

If you remove the control board on the primary Routing Engine, resulting in removal of active fabric planes, the line card takes the local action of disabling the removed planes. If spare planes are available, the line card initiates switchover to spare planes. If an active control board on a backup Routing Engines is removed, the primary Routing Engine disables the removed planes and performs the switchover to spare planes, if available. The software attempts to optimize the duration of connectivity loss by disabling all removed planes. The spare planes are transitioned to the online state one by one.

Fabric self-ping is a mechanism to detect any issues in the fabric data path. Each Packet Forwarding Engine forwards fabric data cells that are destined to itself over all active fabric planes. To transmit the data cell, the Packet Forwarding Engine fabric sends the request cells over an active plane and waits for a grant packet. The destination Packet Forwarding Engine sends a grant packet over the same plane on which the request cell is received. When the grant cell is received, the source Packet Forwarding Engine sends the data cell.

The Packet Forwarding Engine fabric contains the capability to detect grant delays. If grants are not received within a certain period of time, a destination timeout is declared. Destination timeout on a certain plane by a Packet Forwarding Engine on two or more line cards is considered as an indication for plane failures. Even if one Packet Forwarding Engine on a line card flashes an error, the line card is considered to be in error. Destination timeouts are noticed when the Packet Forwarding Engine sends traffic actively because requests are sent only for valid data cells. The software takes an appropriate action based on the destination timeout. For self-ping, a data cell is destined to the source Packet Forwarding Engine only.

Fabric ping failure messages are sent to the fabric manager on the Routing Engine, which collates all of the errors reported by all the line cards and takes a corrective action. For example, a ping failure for all links of the same line card might indicate a problem on the line card. Ping failure for multiple line cards for the same fabric plane might indicate a problem with the fabric.

If the Routing Engine determines that a fabric plane is down, based on the information on errors it receives from the line cards or the Packet Forwarding Engines, over a period of 5 seconds, it indicates a fabric failure. The duration of 5 seconds is the period for which the Routing Engine collates the errors from all of the line cards.

Fabric self-ping packets are periodically sent to check the sanity of the fabric links. Self pings are sent at interval of 500 ms. The destination timeout is also checked in intervals of 500 ms. If two timeouts ocur successively, self ping failure is detected. When a destination timeout is received, the Packet Forwarding Engine fabric stops the sending of packets to the fabric. To examine the link condition again, the software resets the credits to ensure that new requests are sent again. When a self-ping failure occurs, the line card removes the affected plane from sending data to all destinations. This method ensures that self-ping is not attempted to be sent again on the defective plane.

The following guidelines apply to the self-ping capability:

- By default, self pings are not sent on spare fabric planes because spare planes do not carry traffic.
- The size of self-ping packets is large enough to enable the cells to be loaded over all the active fabric planes (MX2020 supports 24 fabric planes and MX10008 supports 12 fabric planes).
- A detection of received self-ping packets is not performed.
- High priority queue is used to enable self-ping to be sent for oversubscription cases.

Detection and Corrective Actions of Line Cards on MX Series Routers

You can configure a line card to be moved to the offline state on an MX-Series routers (such as MX10008, MX10004, MX2020, MX2010, MX2008, MX960, MX480, or MX304, MX240, and so on). Configuring this feature does not affect the system. You can configure this feature without restarting the line card or restarting the system.

The following scenarios can occur when you configure the feature to disable line cards :

-
- If a line card has been brought offline because of fabric errors and this functionality to move the line card to offline state is disabled, the line card is transitioned to the online state automatically.
- If a line card has been brought offline because of fabric errors and this functionality to move the line card to offline state is disabled or configured for some other line card, the line card that was turned offline is transitioned to the online state automatically.
- All the line cards that were brought offline , when you configured this setting, are brought back online when you commit any configuration under the `[edit chassis]` hierarchy level. Similarly, a restart of the chassis daemon or the *Graceful Routing Engine switchover* (GRES) operation also causes the line card that is disabled because of degraded fabric to be moved to the online state.

When a line card is operating with less than the required number of active fabric planes. If a line card is operating with less than four planes, the fabric traffic operates at a reduced bandwidth.

The following conditions can result in reduced operating bandwidth in fabric:

- The fabric control boards go offline as a result of an unintentional, abrupt power shutdown.
- An application-specific integrated circuit (ASIC) error, which causes a plane of a control board to be automatically turned offline.
- Manually bringing the fabric plane or the control board to the offline state.
- Removal of the control board
- Self-ping failure on any plane.
- HSL2 training failure for active plane.
- If a spare fabric plane has CRC errors, and this spare plane is made online, the link with the CRC error is disabled. This mechanism might cause a degradation in fabric in one direction and might cause a null route in the other direction.
- When a self-ping or HSL2 training failure occurs, the fabric plane is disabled for a particular line card and it is online for other line cards. This condition can also cause a null route.

If you need to remove the control board or move a fabric plane to the offline state during a system maintenance, you must enable the functionality to turn the line cards with degraded bandwidth to the offline state (by using the `offline-on-fabric-bandwidth-reduction` statement at the `[edit chassis fpc` *`slot-number`*`]` hierarchy level).

The following corrective actions are performed when a null route or reduced operating bandwidth occurs in the fabric:

- Regardless of whether a spare control board is available or not, self-ping state for each line card is monitored at intervals of 5 seconds at the Routing Engine. Fabric manager determines the presence of spare control boards
- The switch fabric is hosted on the Switch Fabric Boards (SFBs) on MX10008, MX10004, MX2020, MX2010 and MX2000 devices:
  - The MX10008 router has eight slots for the line cards that can support a maximum of 768 100-Gigabit Ethernet ports (4x100), 192 40-Gigabit Ethernet ports, 192 100-Gigabit Ethernet ports, or 192 400-Gigabit Ethernet ports with line card slots 0-7 that combine Packet Forwarding Engine (PFE) and Ethernet interfaces enclosed in a single assembly. MX10008 supports six Switch Fabric Boards (SFBs) There are two models of SFBs: the JNP10008-SF and the JNP10008-SF2. SFBs installed must be of the same model type in a running chassis.

    For details, see No Link Title
  - MX10004 features a compact 7-U modular chassis, line card slots 0-3 silicon line cards (2.4 Tbps, 480 Gbps, and 9.6 Tbps throughput) , with full hardware redundancy. Switch Fabric Boards (SFBs) create the switch fabric for the MX10004. Each SFB has a set of connectors to the line cards and the Routing and Control Board (RCB) to the switch fabric. Three SFBs provide reduced switching

functionality to an MX10004 router. Six SFBs provide full throughput. Each MX10004 SFB has four connectors. Each connector matches up with a line card slot, eliminating the need for a backplane.

For details on fabric plane management, see No Link Title.

- The MX10003 router contains modular routing engines and PFEs. The single PFE performs both ingress and egress packet forwarding. The router provides two dedicated line card slots. The router supports one primary and two redundant Routing and Control Boards (RCBs).

- The MX2020 and MX2010 devices support 8 SFBs. The Mx2020 has 20 dedicated line card slots.The MX2010 router has 10 dedicated line-card slots The host subsystem consists of two Control Boards with Routing Engines (CBREs) and eight Switch Fabric Boards (SFBs). Data packets are transferred across the backplane between the MPCs through the fabric ASICs on the SFBs.

  Switch Fabric Boards (SFBs) provide increased fabric bandwidth per slot. Up to eight SFBs, SFB2s, or

  SFB3s can be installed in an MX2020 or MX2010 router. All switch fabric boards in the chassis must be the same type. Mixed mode is not supported.

- MX960 routers with I-chip or I-chip and Trio-chip-based line cards that contain three control boards.

- MX240 or MX480 routers with I-chip or I-chip and Trio-chip-based line cards that contain two control boards.

- MX960, MX480, or MX240 routers that contain only Trio-based line cards are not considered to contain a spare control board.

If during any such interval of 5 seconds, two line cards indicate a failure for the same plane, a switchover to the spare control board. In this case, the control board that reported errors is turned offline and the spare control board is turned online.

- If a spare control board is available, and if you configure the functionality to disable line cards , self-ping state for each line card is monitored at intervals of 5 seconds at the Routing Engine. The following conditions can occur:

  - During any 5-second interval, if only one line card indicates a failure for a plane, the fabric Manager waits for the next interval. During the subsequent interval, if no other line card indicates a failure for the same plane, switchover of the control board is performed.

  - During any 5-second interval, if multiple line cards show failures for multiple control boards, the fabric manager waits for the next interval. During the subsequent interval, if the same condition remains, all the failing line cards are turned offline even if the spare control board is present.

  - During any 5-second interval, if any line card shows a failure for multiple planes on multiple control boards, the fabric manager waits for the next interval. During the subsequent interval, if the same condition persists, the line card is turned offline even if the spare control board is present.

- If spare planes are not available, the line card is turned offline when it displays a failure for a single plane or multiple planes. The line card is brought offline only if you previously configured the `offline-on-fabric-bandwidth-reduction` statement at the `[edit chassis fpc *slot-number*]` hierarchy level.

Understanding Fabric Fault Handling on T4000 Router

The T4000 router consists of a Switch Interface Board (SIB) with fabric bandwidth double the capacity of the T1600 router. The fabric fault management functionality is similar to that in T1600 routers. This topic describes the fabric fault handling functionality on T4000 routers.

The fabric fault management functionality involves monitoring all high-speed links connected to the fabric and the ones within the fabric core for link failures and link errors.

Action is taken based on the fault and its location. The actions include:

- Reporting link errors in system log files and sending this information to the Routing Engine.
- Reporting link failures at the Flexible Port Concentrator (FPC) or at the SIB and sending this information to the Routing Engine.
- Marking a SIB in `Check` state.
- Moving a SIB into `Fault` state.

The SIB in T4000 routers forms the core of the fabric with 4:1 redundancy—the redundant SIB becomes active when the active SIB becomes nonfunctional, is deactivated, or is removed. The following are the high-level indications of fabric faults that are monitored by Junos OS:

- An SNMP trap is generated whenever a SIB is reported as `Check` or `Fault`.
- `show chassis alarms`—Indicates that a SIB is in `Check` or `Fault` state.
- `show chassis sibs`—Indicates that a SIB is in `Check` or `Fault` state or that a SIB is in `Offline` state when the SIB initializes (this occurs when the SIB does not power on fully).
- `show chassis fabric fpcs`—Indicates whether any fabric links are in error on the FPCs' side.
- `show chassis fabric sibs`—Indicates whether any fabric links are in error on the SIBs' side.

- The /var/log/messages system log messages file at the Routing Engine has error messages with the prefix CHASSISD_FM_ERROR.
- The SIBs display the FAIL LED.

NOTE: The fabric planes in the chassis determine whether the chassis is a T640 router, a T1600 router, or a T4000 router. Power entry modules (PEMs), FPCs, or fan trays do not determine chassis personality. Alarms are raised if the old PEMs or fan trays are present in a T4000 chassis. You can identify a router based on its fabric planes:

- If all planes present are F16-based SIBs, the chassis is a T640 chassis.
- If all planes present are SF-based SIBs, the chassis is a T1600 chassis.
- If all planes present are XF-based SIBs, the chassis is a T4000 chassis.

Note that mixing of fabric planes is not a supported configuration except during upgrade. You can change the personality of a chassis without a reboot by changing all the fabric planes and by issuing the `set chassis fabric upgrade-mode` CLI command to check the personality. If you do not issue the `set chassis fabric upgrade-mode` CLI command, the personality does not change until the next boot.

In T4000 routers, you come across the following faults:

- Board-level faults—These faults occur during initialization or during runtime. Power failure during board initialization, high-speed links transmit error, and polled I/O error during runtime are some examples of board-level faults.
- Link-level faults—These faults occur during initialization or during runtime. Link training failure at initialization time (failure of the data plane links between an FPC and a SIB to be trained when the FPC or SIB is initialized), error detected on the channel between the SIB and a Packet Forwarding Engine, cyclic redundancy check (CRC) errors detected at runtime, and Packet Forwarding Engine destination errors are types of link-level faults.
- Faults based on environmental conditions—These faults occur during runtime. Sudden removal of an FPC or a SIB might result in an operator error. When a SIB becomes too hot or when SIB voltages are beyond thresholds, the errors generated are classified into environmental errors.

You can implement one of the following options to handle the faults:

- Log the error and raise an alarm.
- Switch over to the spare plane, if available.
- Continue with a reduced number of parts of a plane.
- Continue with a reduced number of usable planes.

- Use polling-based fault handling.
- Monitor high-speed link errors and manually bring the link down to a suitable threshold.

The polled I/O errors and the link errors are monitored every 500 milliseconds, and the board exhaust temperature and board voltages are monitored every 10 seconds.

Understanding Fabric Fault Handling on PTX5000 Packet Transport Router

IN THIS SECTION

Starting with Junos OS Release 14.1, the PTX5000 Packet Transport Router supports nine Switch Interface Boards (SIBs). Each FPC2-PTX-P1A FPC supports 1Tb per slot capacity, thereby resulting in a fabric bandwidth of 16 terabits per second (Tbps), full-duplex (8 Tbps of any-to-any, nonblocking, half-duplex) switching.

The fabric fault management functionality involves monitoring all high-speed links connected to the fabric and the ones within the fabric core for link failures and link errors.

The faults that occur in a PTX5000 can be broadly categorized into:

- Board faults—Faults that arise in a SIB or in an Flexible Port Concentrator (FPC) during initialization or during runtime, including issues that arise when a router component is accessing the SIB or FPC or issues that arise out of midplane failures.
- Link faults—Faults that occur on high-level links in a router during initialization or during runtime.
- Faults due to environmental conditions—Faults that occur because of overvoltage or over-temperature; faults that occur because of an operator mishandling a SIB or an FPC, and so on.

The router takes action on the basis of the fault category and the fault location. The actions include:

- Reporting link errors in system log files and sending this information to the Routing Engine.
- Displaying the link errors when you run one of the operational commands listed in Table 2 on page 43:

Table 2: List of Operational Mode Commands

Operational mode command	Description
`show chassis sibs`	Displays Switch Interface Boards (SIBs) status information.
`show chassis fabric fpcs <slot number>`	Displays the fabric state of the specified FPC slot. If no slot number is provided, it displays the status of all FPCs.
`show chassis fabric sibs <slot number>`	Displays the state of the electrical switch fabric link between the SIBs and the FPCs.
`show chassis fabric reachability <detail>`	Displays the current state of fabric destination reachability.
`show chassis fabric unreachable-destinations`	Displays the list of destinations that have transitioned from a reachable state to an unreachable state.
`show pfe statistics error`	Displays Packet Forwarding Engine error statistics.
`show chassis fabric topology <sib_slot>`	Displays the input-output link topology.
`show chassis fabric summary`	Displays the state of all fabric planes and the elapsed uptime.

- Reporting link failures at the FPC level or at the SIB level and sending this information to the Routing Engine.
- Reporting link error information in the `show chassis alarms` operational command.
- Moving a SIB into *fault* state.

The following sections explain fabric fault handling functionality on the PTX5000:

SIB-Level Faults

The following sections give a brief overview on the types of faults that occur on a SIB and how to handle them:

Types of Faults That Occur on a SIB

Board faults and link faults occur on a SIB during initialization and during runtime. Some faults occur because of environmental conditions such as overvoltage or over-temperature, or when an operator mishandles the SIB.

> **NOTE**: Run the operational mode commands listed in Table 2 on page 43 to detect faults.

During SIB initialization and runtime, the following faults might occur:

- Board faults, such as failure of SIBs to power up, ASICs reset failure, Switch Processor Mezzanine Board (SPMB) polled I/O access failure to ASICs, board component failures such as PIC failures, or router component access failures.
- Link faults such as high-level link errors that occur during *link training*.
- Faults that occur because of environmental conditions or because of mishandling of the SIB by the operator.

Handling SIB-Level Faults

The following list illustrates how the router handles a fault that occurs on a SIB during initialization, during runtime, because of environmental conditions, and because of mishandling of the SIB by the operator:

- To handle a board fault on a SIB during initialization, the chassis daemon (*chassisd*) marks the SIB to be in *fault* state. After the SIB is marked as faulty, no operation occurs on this SIB.
- To handle a board fault on a SIB during runtime, chassisd logs an error in the system log file, raises an alarm indication error type, and marks the SIB as faulty. After the SIB is marked as faulty, no operation occurs on this SIB.
- To handle a link fault on a SIB during runtime, when a link error comes up during link training, chassisd informs the FPC corresponding to the link on which the error occurred to disable the links to the affected SIB. The chassisd then sends an error message to all the other FPCs in the router to stop using the failed SIB link and a link error alarm is generated. Note that when more than one FPC report errors for a given SIB, the SIB is disabled for all FPCs and no traffic is sent by the Packet Forwarding Engine through the affected SIB.
- To handle a link fault on a SIB during runtime, chassisd marks the SIB as faulty and specifies a reason for the error, and the SIB is disabled.

- In case of an environmental fault—overvoltage or over-temperature—the SIB is immediately taken offline. Note that an error is logged periodically as the temperature or voltage rises, and the SIB is taken offline when it crosses a certain threshold voltage or temperature.
- When a SIB is abruptly removed or dislodged, all the affected Packet Forwarding Engines stop using that plane to reach other Packet Forwarding Engines in the router.

FPC-Level Faults

The following sections give a brief overview of the types of faults that occur on an FPC and how to handle them:

Types of Faults That Occur on an FPC

Board faults and link faults occur on an FPC during initialization and during runtime. Some faults also occur because of environmental conditions such as overvoltage, over-temperature, or when the operator mishandles the FPC.

> **NOTE**: Run the operational commands listed in Table 2 on page 43 to detect faults.

During FPC initialization and runtime, the following faults might occur:

- Board faults such as failure of FPCs to power up, failure of ASICs to come out of reset phase, PMB polled I/O access failure to ASICs, board component failures such as PIC failure, or router component access failures.
- Link faults such as high-level link errors that occur during link training.
- Faults that occur because of environmental conditions or because of mishandling of an FPC by the operator.

Handling FPC-Level Faults

The following list illustrates how the router handles a fault that occurs on an FPC during initialization, during runtime, because of environmental conditions, and because of mishandling of the FPC by the operator:

- To handle a board fault on an FPC during initialization, chassisd marks the FPC to be in *fault* state. After the SIB is marked as faulty, no operation occurs on this FPC.
- To handle a board fault on an FPC during runtime, chassisd logs an error in the system log file, raises an alarm indication error type, and marks the FPC as faulty. After the FPC is marked as faulty, no operation occurs on this FPC.

- To handle onboard link errors on an FPC during initialization or during runtime, the FPC is taken down and all the affected Packet Forwarding Engines stop using that plane to reach other Packet Forwarding Engines in the router.

> **NOTE**: No planes are taken down during initialization because the link training process for the fabric is not yet complete.
>
> Onboard link errors during runtime are resolved on the basis of current configuration; either the FPC is rebooted or the error is logged and the FPC continues with initialization.

- In case of an environmental fault—over voltage or over-temperature—the FPC is immediately taken offline. Note that an error is logged periodically as the temperature or voltage rises, and the FPC is taken offline when it crosses a certain threshold voltage or temperature.

- When an FPC is abruptly removed or dislodged, all the other Packet Forwarding Engines stop sending traffic to the Packet Forwarding Engines in this FPC.

Understanding Fabric Fault Handling on Enhanced Switch Fabric Board (SFB2)

The MX2000 line of routers support Switch Fabric Boards (SFBs) and enhanced SFBs (SFB2s) but not both at the same time. The SFB and SFB2 host three fabric planes each. So, the chassis supports a total of 24 planes. Junos OS Release 15.1F6 and 16.1R1 support fabric fault handling for each plane in both SFB and SFB2. In earlier releases, fabric fault handling is supported for each SFB, not for each plane.

Table 3 on page 46 lists the differences between fabric fault handling per plane and per SFB.

Table 3: SFB Versus SFB2 Fabric Fault Handling

SFB Level (SFB)	Plane Level (SFB and SFB2)
Cyclic redundancy check(CRC) errors on any link on the SFB are indicated on the SFB.	CRC errors on any link on the SFB or SFB2 are indicated on the plane.
On encountering destination errors, the line card isolates the SFB (all 3 planes).	On encountering destination errors, the line card isolates the corresponding plane. Other planes continue to operate.

Fabric fault handling per-plane provides the following benefits:

- Increased granularity, which helps identify, isolate, and repair faults.
- Alarms and log messages provide fault information per plane instead of per SFB, which makes debugging easier.
- If an SFB has a single faulty plane, the other two planes can continue to operate. There is no need to take the entire SFB offline.
- In case of transient errors, while repairing you can isolate a single plane instead of isolating the bouncing the SFB.

To view fabric fault handling information for all 24 planes, use the `extended` option with the existing fabric commands.

Managing Bandwidth Degradation

Certain errors result in packets being dropped by a system without notification. Other connected systems continue to forward traffic to the affected system, impacting network performance. A severely degraded fabric plane can be one of the reasons here.

By default, Juniper Networks routers attempt to start healing from such situations when the system detects issues with Packet Forwarding Engines. If the healing fails, the system turns off the interfaces, thereby preventing further escalations.

On Junos OS, you can use the configuration statement `bandwidth-degradation` at the `[edit chassis fpc slot-numberfabric]` hierarchy to detect and respond to fabric plane degradation in ways you deem fit. You can configure the router to specify which healing actions the router should take once such a condition is detected. You can also use the optional statement `blackhole-action` to determine how the line card responds to a 100 percent fabric degradation scenario. This command is optional and overrides the default fabric hardening procedures.

NOTE: The `bandwidth-degradation` command and the `offline-on-fabric-bandwidth-reduction` statements are mutually exclusive. If both commands are configured, an error is issued during the commit check.

The `bandwidth-degradation` statement is configured with a percentage and an action. The `percent-age` value can range from 1 to 99, and it represents the percentage of fabric degradation needed to trigger a response from the line card. The `action` attribute determines the type of response the line card performs once fabric degradation reaches the configured percentage.

The statement is only configured with an `action` attribute, which triggers when the percentage of fabric degradation reaches 100 percent.

The following actions can be applied to either configuration statement:

- `log-only`: A message gets logged in the chassisd and message files when the fabric degradation threshold is reached. No other actions are taken.
- `restart`: The line card with a degraded fabric plane is restarted once the threshold is reached.
- `offline`: The line card with a degraded fabric plane is taken offline once the threshold is reached. The line card requires manual intervention to be brought back online. This is the default action if no action attribute configured.
- `restart-then-offline`: The line card with a degraded fabric plane is restarted once the threshold is reached, and if fabric plane degradation is detected again within 10 minutes, the line card is taken offline. The line card requires manual intervention to be brought back online.

NOTE: This feature is available in the Junos OS Release 15.1R1.

Fabric Hardening and Recovery on PTX10001-36MR, PTX10004, PTX10008, and PTX100016 with PTX10K-LC1202-36MR Line Card

IN THIS SECTION

PTX10001-36MR, PTX10004, PTX10008, and PTX100016 routers support fabric hardening. Fabric hardening is a resiliency feature to detect fabric blackholing and attempt automatic recovery process to restore the Packet Forwarding Engines from blackhole condition.

We've enabled fabric hardening by default. When the system detects any unreachable Packet Forwarding Engine destination, this feature attempts automatic fabric connectivity restoration.

If restoration fails, the system turns off the interfaces to limit the blackholing and trigger alarm to indicate the unreachable Packet Forwarding Engine destinations. However, instead of turning off the interfaces, user can configure Packet Forwarding Engine offline by using `set chassis fabric event reachability-fault actions recovery-failure pfe-offline` statement at the `[set chassis fabric event]` hierarchy level.

Packet Forwarding Engine destinations can become unreachable for the following reasons:

- Complete self-blackhole- Complete connectivity loss occurs on all fabric planes.
- Complete peer-blackhole- Two Packet Forwarding Engines can reach the fabric but not each other.

You can configure a router to trigger fabric recovery when the router detects degradation in fabric bandwidth by using `degraded` statement at the `[edit chassis fabric event reachability-fault]` hierarchy level. The degradation statement is configured with a percentage value that can range from 1 to 99. The percentage value represents the error threshold for fabric bandwidth degradation and the router starts the recovery once the threshold is reached.

When the degraded error threshold is configured, the router can also attempt fabric recovery for the following reasons:

- Self degrdation- Degraded fabric condition in a Packet Forwarding Engine destination.
- Peer degradation- Degraded fabric condition between two Packet Forwarding Engines.

The fabric recovery process involves one or more of the following phases:

- SIB restart phase: If Packet Forwarding Engine destinations across multiple line cards have fabric connectivity failures on planes, then the router attempts to resolve the issue by restarting the SIBs. If multiple SIBs require a restart, the router restarts the SIBs one by one.
- FPC restart phase: The router attempts automatic recovery by restarting the FPCs for the following scenarios:
  - All Packet Forwarding Engine destinations having complete or partial blackhole conditions are in a single FPC.
  - If Packet Forwarding Engine destinations with complete or partial blackhole conditions occur across different FPCs, but none of the Packet Forwarding Engines share common plane of failure.
  - The attempt of SIB restart phase failed to recover Packet Forwarding Engines.

  You can disable restarting of FPCs to limit recovery actions from a degraded fabric condition. To disable restarting of FPCs, use the `set chassis fabric event reachability-fault actions fpc-restart-disable` statement at the `[set chassis fabric event]` hierarchy level.
- Packet Forwarding Engine offline phase: Because previous attempts of recovery phases failed or recovery action disabled in the configuration, the router turns off the interfaces to limit the blackholing by default. However, instead of turning off the interfaces, user can configure Packet Forwarding Engine offline by using `set chassis fabric event reachability-fault actions recovery-failure pfe-offline` statement at the `[set chassis fabric event]` hierarchy level.

If the router has only Packet Forwarding Engines with peer blackhole or peer degradation condition, then the router attempts recovery through link autoheal by restarting fabric links on the planes.

Benefits

- Attempts automatic recovery process to recover the Packet Forwarding Engines from degraded fabric conditions to minimize traffic loss.
- Raise alarms that provide fault information to indicate the unreachable Packet Forwarding Engine destinations, if the recovery fails.

Disabling Line Card Restart to Limit Recovery Actions from Degraded Fabric Conditions

You can disable line card restarts to limit recovery actions from a degraded fabric condition. On T640 and T1600 routers, only the fabric plane is restarted. On PTX Series routers, only the Switch Interface Boards (SIBs) are restarted. To disable the restarting of line cards, use the `action-fpc-restart-disable` statement at the `[edit chassis fabric degraded]` hierarchy level:

```
[edit chassis fabric]
degraded
```

Whenever a line card restart is disabled, an alarm is raised when there are unreachable destinations present in the router, and you must restart the line cards manually.

To ensure that both the fabric planes (T640 and T1600 routers) or the SIBs (PTX Series routers) and the line cards are restarted during the recovery process, do not configure the `action-fpc-restart-disable` statement at the `[edit chassis fabric degraded]` hierarchy level.

Disabling an FPC with Degraded Fabric Bandwidth

You can bring an FPC with degraded fabric bandwidth offline to avoid causing a null route in the chassis for an extended time. To configure the option to disable an FPC with degraded bandwidth, use the `offline-on-fabric-bandwidth-reduction` statement at the `[edit chassis fpc slot-number]` hierarchy level:

```
[edit chassis]
fpc slot-number {
    offline-on-fabric-bandwidth-reduction;
}
```

The fabric manager checks the number of current active planes periodically. If the number of active planes is lower than the required number of active planes for a particular router, the system waits 10 seconds before it takes any corrective action. If the reduced bandwidth condition persists for an FPC and if this feature has been configured for the FPC, the system brings the FPC offline.

Error Handling by Fabric OAM

Fabric Operation, Administration, Maintenance (OAM) helps in detecting failures in fabric paths. Fabric OAM validates the fabric connectivity before sending traffic on a fabric plane whenever a new fabric path is brought up for a PFE. If a failure is detected, the software reports the fault and avoids using that fabric plane for that PFE. This feature works by sending a very low packets per second (PPS) self-destined OAM traffic over each of the available fabric planes and detecting any loss of traffic at the end points (fabric self-ping check).

NOTE:

- In Junos OS Evolved Release 20.4R1, the fabric OAM feature is enabled by default. You can disable the feature by using the CLI command `set chassis fabric oam detection-disable`.
- In Junos OS Evolved Releases 20.4R2 and21.1R1, the fabric OAM feature is disabled by default.
- In Junos OS Evolved Release 22.1R1, the runtime fabric OAM feature is enabled by default. You can disable the feature by using the CLI command `edit chassis fabric oam runtime-disable`. The runtime fabric OAM feature is supported on PTX10004, PTX10008, and PTX10016 routers.

The Fabric OAM checks are done at boot time. The failed paths are disabled. The system does not do any recovery action. However, you can try to recover the affected fabric planes by restarting the SIBs. The recovery steps depend on the nature of the failure.

A fabric plane represents an independent bidirectional path between a PFE and fabric ASIC. Runtime Fabric OAM periodically checks fabric connectivity and helps detect and report failures in fabric planes during system runtime. Runtime Fabric OAM detects the fabric reachability of each PFE.

When the same fabric planes fail on a single or multiple FPCs, restart the SIB containing the failed planes, using the following commands:

```
user@host> request chassis sib slot slot-number offline
```

```
user@host> request chassis sib slot slot-number online
```

When random fabric planes fail on multiple FPCs, the fault cannot be isolated to a specific FPC or SIB. However, you can try to recover the planes by restarting the SIBs that contain the affected planes in a sequential manner.

For each error detected by the fabric OAM feature, a syslog is generated. The following is an example:

```
Oct 29 23:02:46  router-dvi resiliencyd[12921]: Error: /fpc/0/fabspoked-pfe/0/cm/0/pfe/0/
fabric_link_foam_fault (0x410009), scope: board, category: internal, severity: major, module:
fab-pfe@0, type: fabric link foam fault
```

The following syslog message indicates that a fabric OAM-related error was cleared.

```
Oct 29 23:25:14  router-dvi resiliencyd[12921]: Performing action clear-cmalarm for error /fpc/0/
fabspoked-pfe/0/cm/0/pfe/0/fabric_link_foam_fault (0x410009) in module: fab-pfe@0 with scope:
board category: internal level: major
```

Also, you can use the CLI commands `show system errors active detail` and `show system alarms` to view the Fabric OAM-related errors.

```
user@router> show system alarms
20 alarms currently active
Alarm time               Class  Description
2020-08-20 10:32:02 UTC  Major  FPC 0 Ideeprom read failure
2020-08-20 10:58:07 UTC  Major  FPC 0 Self_FOAM fault detected
[...Output truncated...]
```

```
user@router> show system alarms
14 alarms currently active
Alarm time               Class  Description
2022-02-15 23:45:28 PST  Minor  FPC 1 Volt Sensor Fail
2022-02-16 00:02:03 PST  Major  FPC 1 Self_Fabric OAM Runtime fault detected
2022-02-15 23:43:04 PST  Minor  FPC 1 Secure boot disabled or not enforced
2022-02-15 23:55:50 PST  Minor  FPC 3 Secure boot disabled or not enforced
[...Output truncated...]
```

The following output shows details for both single fabric plane failure (on Packet Forwarding Engine 0) and all fabric planes failure (on Packet Forwarding Engine 1).

```
user@router> show system errors active detail
System Active Errors Detail Information
FPC 0
----------------------------------------------------------------
Error Name           : fabric_down_condition_on_pfe
Identifier           : /fpc/0/fabricHub/0/cm/0/fabrichub/1/fabric_down_condition_on_pfe
Description          : fabric_down_condition_on_pfe
State                : enabled
Scope                : pfe
Category             : functional
Level                : major
Threshold            : 1
Error limit          : 0
Occur count          : 3
Clear count          : 2
Last occurred(ms ago) : 103158
System Active Errors Detail Information
FPC 0
----------------------------------------------------------------
Error Name           : fabric_link_foam_fault
Identifier           : /fpc/0/fabspoked-pfe/0/cm/0/pfe/0/fabric_link_foam_fault
Description          : fabric link foam fault
State                : enabled
Scope                : board
Category             : internal
Level                : major
Threshold            : 1
Error limit          : 100
Occur count          : 2
Clear count          : 0
Last occurred(ms ago) : 113277
System Active Errors Detail Information
FPC 0
----------------------------------------------------------------
Error Name           : fabric_link_foam_fault
Identifier           : /fpc/0/fabspoked-pfe/0/cm/0/pfe/1/fabric_link_foam_fault
Description          : fabric link foam fault
State                : enabled
Scope                : board
```

```
Category              : internal
Level                 : major
Threshold             : 1
Error limit           : 100
Occur count           : 12
Clear count           : 0
Last occurred(ms ago) : 103267
System Active Errors Detail Information
RE 0
----------------------------------------------------------------
Error Name            : fpga_min_supported_fw_ver_mismatch
Identifier            : /re/0/hwdre/0/cm/0/fpga_fw_events/UBAM FPGA/
fpga_min_supported_fw_ver_mismatch
Description           : firmware_version_lower_than_minimum_expected
State                 : enabled
Scope                 : board
Category              : functional
Level                 : minor
Threshold             : 10
Error limit           : 1
Occur count           : 1
Clear count           : 0
Last occurred(ms ago) : 68886367

FPC 1
----------------------------------------------------------------
Error Name            : fabric_link_self_fabric_oam_runtime_fault
Identifier            : /fpc/1/fabspoked-pfe/0/cm/0/pfe/0/
fabric_link_self_fabric_oam_runtime_fault
Description           : fabric link self fabric oam runtime fault
State                 : enabled
Scope                 : board
Category              : internal
Level                 : major
Threshold             : 1
Error limit           : 36
Occur count           : 1
Clear count           : 0
Last occurred(ms ago) : 2022-02-16 00:02:03 PST (448108 ms ago) System Active Errors Detail
Information
```

You can use the CLI command `show chassis fabric fpcs` to view the fabric OAM self-ping state of each fabric plane.

```
user@router> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
    PFE #0
SIB0_Asic0_Fcore0 (plane 0)  Plane Disabled, Links ok Fabric OAM failed
SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok Fabric OAM success
SIB0_Asic0_Fcore0 (plane 2)  Plane Enabled, Links ok Fabric OAM success
SIB0_Asic0_Fcore0 (plane 3)  Plane Enabled, Links ok Fabric OAM success
SIB0_Asic0_Fcore0 (plane 4)  Plane Enabled, Links ok Fabric OAM success
SIB0_Asic0_Fcore0 (plane 5)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 8)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 9)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 10)  Plane Enabled, Links ok Fabric OAM success
SIB1_Asic0_Fcore0 (plane 11)  Plane Enabled, Links ok Fabric OAM success
    PFE #1
SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok Fabric OAM success
SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok Fabric OAM success
```

```
user@router> show chassis fabric fpcs
Fabric management FPC state:
FPC #1
    PFE #0
    SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic0_Fcore0 (plane 1)  Plane Disabled, Links ok Fabric OAM Runtime failed
    SIB0_Asic1_Fcore0 (plane 2)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic1_Fcore0 (plane 3)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic2_Fcore0 (plane 4)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB0_Asic2_Fcore0 (plane 5)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic1_Fcore0 (plane 8)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic1_Fcore0 (plane 9)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic2_Fcore0 (plane 10)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB1_Asic2_Fcore0 (plane 11)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic0_Fcore0 (plane 12)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic0_Fcore0 (plane 13)  Plane Enabled, Links ok Fabric OAM Runtime success
```

```
    SIB2_Asic1_Fcore0 (plane 14)  Plane Enabled, Links ok Fabric OAM Runtime success
    SIB2_Asic1_Fcore0 (plane 15)  Plane Enabled, Links ok Fabric OAM Runtime success
```

The `show chassis fabric fpcs` command displays the following output when the fabric OAM feature is disabled:

```
user@router> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
    PFE #0
        SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 2)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 3)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 4)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 5)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 8)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 9)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 10)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 11)  Plane Enabled, Links ok
    PFE #1
        SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 2)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 3)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 4)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 5)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 8)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 9)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 10)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 11)  Plane Enabled, Links ok
    PFE #2
        SIB0_Asic0_Fcore0 (plane 0)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 1)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 2)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 3)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 4)  Plane Enabled, Links ok
        SIB0_Asic0_Fcore0 (plane 5)  Plane Enabled, Links ok
```

```
        SIB1_Asic0_Fcore0 (plane 6)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 7)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 8)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 9)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 10)  Plane Enabled, Links ok
        SIB1_Asic0_Fcore0 (plane 11)  Plane Enabled, Links ok
    PFE #3
```

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
14.2R6	Starting in Junos OS Release 14.2R6, if a SIB becomes offline because of extreme conditions such as high voltage or high temperature, then as part of the recovery process, the router does not restart the fabric plane for that SIB.
14.2R6	Starting in Junos OS Release 14.2R6, you can manage fabric degradation in single-chassis systems better by incorporating fabric self-ping and Packet Forwarding Engine liveness mechanisms.
14.1	Starting with Junos OS Release 14.1, the PTX5000 Packet Transport Router supports nine Switch Interface Boards (SIBs).
13.3	Starting with Junos OS Release 13.3, you can use PTX Series routers to configure Packet Forwarding Engine (PFE)-related error levels and the actions to perform when a specified threshold is reached.

RELATED DOCUMENTATION

Fabric Grant Bypass | 75
Fabric Plane Management | 58
Troubleshooting the T4000 SIBs
Troubleshooting the T4000 FPCs
show chassis alarms
show chassis fabric fpcs
show chassis fabric destinations
show chassis fabric sibs
show chassis sibs
PTX5000 Description

FPCs Supported on the PTX5000
PTX5000 Switch Interface Board Description

Fabric Plane Management

IN THIS SECTION

Configuring Fabric Redundancy Mode for Active Control Boards on MX Series Routers

An MX960 router can support three Enhanced Switch Control Boards (SCBE2s or SCBEs)—two planes on each SCB and make up a total of six fabric planes. MX240 and MX480 routers can support up to two SCBE2s or SCBEs—four fabric planes on each SCBE make up a total of eight planes. However, the MX240 and MX480 routers have only six active planes. The remaining two are redundant.

MX2020 routers support 8 Switch Fabric Boards (SFBs) or 24 fabric planes. The MX2020 router has 20 dedicated line card slots. The host subsystem consists of two Control Boards with Routing Engines (CBREs). The MX2020 chassis provides redundancy and resiliency. All major hardware components, including the power system, cooling system, control board, and switch fabrics, are fully redundant.

MX10004 supports six SFBs. Each SFB with the switch fabric is connected to the line cards and the Routing and Control Board (RCB). Three SFBs provide reduced switching functionality to an MX10004 router. Six SFBs provide full throughput. Each MX10004 SFB has four connectors. Each connector matches up with a line card slot, eliminating the need for a backplane. The MX10004 power system and the Routing Control Board (RCB) provide redundancy and resiliency.

The MX2010 routers support 8 Switch Fabric Boards (SFBs) and 2 control boards. The MX2010 router provides redundancy and resiliency. All major hardware components, including the power system, cooling system, control board, and switch fabrics, are fully redundant.

An MX10008 devices has six Switch Fabric Boards (SFBs). MX10K-LC2101 has six Packet Forwarding Engines (PFE). Each PFE has 24 connections to the fabric (24 planes, or 4 connections per SFB).

The MX10008 has two models of SFBs: the JNP10008-SF and the JNP10008-SF2. SFBs installed must be of the same model type in a running chassis. On both SF and SF2 models, the SFB has eight connectors that connect to one of the eight line cards.

NOTE: The MPC7E-MRATE and MPC7E-10G MPCs are supported only on MX-SCBE2.

You can configure the active control board to be in redundancy mode or in increased fabric bandwidth mode. You can enable increased fabric bandwidth of active control boards for optimal and efficient performance and traffic handling by configuring the active control boards to be in redundancy mode. To configure redundancy mode for the active control board, use the `redundancy-mode redundant` statement at the `[edit chassis fabric]` hierarchy level:

When you configure this option, all the FPCs use 4 fabric planes as active planes, regardless of the type of the FPC.

To configure increased bandwidth mode for the active control board, use the `redundancy-mode increased-bandwidth` statement at the `[edit chassis fabric]` hierarchy level:

In increased fabric bandwidth mode, MX Series routers will use 6 active planes. MX240 and MX480 routers will also use 2 spare planes in addition to the 6 active planes.

Increased fabric bandwidth mode is enabled by default on MX routers with Switch Control Board (SCB). On MX routers with Enhanced SCB—SCBE, regardless of the type of MPC or DPC installed on it, redundancy mode is enabled by default.

Configuring this feature does not affect the system. You can configure this feature without restarting the FPC or restarting the system.

See also: MX-Series Switch Control Board (SCB) Description

Example: Configuring Fabric Redundancy Mode

IN THIS SECTION

Requirements for Configuration of the Fabric Redundancy Mode

This example uses the following hardware and software components:

- Junos OS Release 12.3 R2 or later for MX Series routers
- A single MX480 router with MPC4E

Overview

This example provides information about configuring the fabric redundancy mode on an MX480 router with MPC4E. You can configure the MPC4E to function in redundant fabric mode or increased bandwidth mode. If you do not configure the mode, the MPC4E, by default, functions in redundant fabric mode. In redundant fabric mode, the number of active fabric planes is 4. If you configure the MPC4E to function in increased bandwidth mode, the number of active fabric planes increases to 6.

See also: 32x10GE MPC4E and 2x100GE + 8x10GE MPC4E.

Configuring Increased Bandwidth Mode

IN THIS SECTION

Procedure

Step-by-Step Procedure

In this example, you configure increased bandwidth mode on an MX480 router with MPC4E. The existing fabric mode on the MX480 router is redundant fabric mode. To configure the fabric mode, perform the following tasks:

1. Verify the existing fabric mode of the router by using the `show chassis fabric mode` command.

```
user@host > show chassis fabric mode
Fabric Operating Mode :
          Redundant Fabric
```

2. View the number of active fabric planes by using the `show chassis fabric summary` command.

```
user@host > show chassis fabric summary
Plane   State    Uptime
 0      Online   2 hours, 58 minutes, 22 seconds
 1      Online   6 seconds
 2      Online   32 seconds
 3      Online   2 hours, 58 minutes, 23 seconds
 4      Spare    31 seconds
 5      Spare    21 seconds
 6      Spare    18 seconds
 7      Spare    9 seconds

For FPC slots with MPC Type 4, Type 5, or MCC:
   Fabric planes 1 and 5, 3 and 7 use shared physical links.
   Those slots may run in a reduced bandwidth in case both
   plane 1 and 5, or both 3 and 7 are active.
```

   Type 4 and Type 5 MPCs refer to MPC 4 and MPC5 line cards, respectively.

3. In configuration mode, go to the `[edit chassis]` hierarchy level and set the fabric mode to `increased-bandwidth` as follows:

```
[edit chassis]
user@ host #set fabric redundancy-mode increased-bandwidth
```

Results

In `redundant fabric` mode, the number of active fabric planes is 4 while the number of spare planes is also 4. In `increased-bandwidth` mode, the number of active planes is 6 while the number of spare planes is 2.

NOTE: Fabric planes 1 and 5 and fabric planes 3 and 7 use shared physical links. So, among fabric planes 1 and 5, only one plane can be active. Similarly, among fabric planes 3 and 7, only one plane can be active.

Verification

IN THIS SECTION

To verify that the fabric mode of the MX480 router with MPC4E, perform the following tasks:

Verifying the Fabric Redundancy Mode of the Router

Purpose

To verify that the fabric redundancy mode of the MX480 router with MPC4E has been modified to `increased-bandwidth`.

Action

To view the fabric mode of the router, use the `show chassis fabric mode` command.

```
user@host > show chassis fabric mode
Fabric redundancy mode: Increased Bandwidth
```

Meaning

The MX480 router with MPC4E is functioning in increased bandwidth mode.

Verifying the Number of Active Fabric Planes

Purpose

To verify that the number of active fabric planes is 6.

Action

To view the number of active fabric planes, use the `show chassis fabric summary` command.

```
user@host > show chassis fabric summary
Plane   State    Uptime
 0      Online   2 hours, 55 minutes, 49 seconds
 1      Online   2 hours, 55 minutes, 25 seconds
 2      Online   2 hours, 58 minutes, 48 seconds
 3      Online   2 hours, 55 minutes, 50 seconds
 4      Online   2 hours, 55 minutes, 48 seconds
 5      Spare    2 hours, 55 minutes, 40 seconds
 6      Online   2 hours, 55 minutes, 37 seconds
 7      Spare    2 hours, 55 minutes, 29 seconds

For FPC slots with MPC Type 4, Type 5, or MCC:
   Fabric planes 1 and 5, 3 and 7 use shared physical links.
   Those slots may run in a reduced bandwidth in case both
   plane 1 and 5, or both 3 and 7 are active.
```

Type 4 and Type 5 MPCs refer to MPC 4 and MPC5 line cards, respectively.

Meaning

Number of active planes on the MX480 router with MPC4E is 6 (0, 1, 2, 3, 4, and 6) while the number of spare planes is 2.

Fabric Plane Management on AS MLC Modular Carrier Card

The Application Services Modular Line Card (AS MLC) provides high application throughput and storage space, and is designed to run services on the MX240, MX480, and MX960 routers. The AS MLC consists of the following components:

- Application Services Modular Carrier Card (AS MCC)

- Application Services Modular Processing Card (AS MXC)
- Application Services Modular Storage Card (AS MSC)

The AS MCC plugs into the chassis and provides the fabric interface.

An MX960 router can support three Switch Control Boards (SCBs) or six fabric planes. The AS MCC supports six fabric planes. An MX240 or MX480 router can support upto two SCBs or two fabric planes. The AS MCC at any time can provide connectivity to only six of the eight fabric planes. Fabric planes 1 and 5, and 3 and 7 use shared physical links. So between fabric planes 1 and 5 only one plane can be active. Similarly between fabric planes 3 and 7, only one plane can be active.

This behavior impacts the output of fabric-related monitoring commands on MX240 and MX480 routers with AS MCCs.

The `show chassis fpc pic-status` command displays the output for an MX480 router with an AS MCC:

```
user@host>show chassis fpc pic-status
Slot 2   Online       MPC Type 1 3D Q
  Slot 1   Online       AS-MCC
  PIC 0  Online       AS-MSC
  PIC 2  Online       AS-MXC
Slot 4   Offline      MPC 3D 16x 10GE
Slot 5   Offline      AS-MCC
```

In the `show chassis fpc pic-status` command output, **Slot 1 and 5** are AS MCC, **PIC 0** is the AS MSC, and **PIC 2** is the AS MXC.

The `show chassis fabric fpcs` command displays the output on an MX480 router with an AS MCC.

```
user@hostshow chassis fabric fpcs
FPC 2
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 4
  PFE #0
      Plane 0: Plane enabled
```

```
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
        Plane 5: Links ok
        Plane 6: Links ok
        Plane 7: Links ok
    PFE #2
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Links ok
        Plane 5: Links ok
        Plane 6: Links ok
        Plane 7: Links ok
  FPC 5
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Unused
        Plane 6: Plane enabled
        Plane 7: Unused
```

In the `show chassis fabric fpcs` command output, **FPC 5** is the AS MCC.

The `show chassis fabric plane` command displays the output on an MX480 router with an AS MCC.

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
```

```
Plane 1
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
Plane 4
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
      FPC 5
          PFE 0 :Links ok
Plane 5
  Plane state: ACTIVE
      FPC 2
          PFE 0 :Links ok
      FPC 4
          PFE 0 :Links ok
          PFE 2 :Links ok
```

```
        FPC 5
            PFE 0 :Unused
  Plane 6
    Plane state: ACTIVE
        FPC 2
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
            PFE 2 :Links ok
        FPC 5
            PFE 0 :Links ok
  Plane 7
    Plane state: ACTIVE
        FPC 2
            PFE 0 :Links ok
        FPC 4
            PFE 0 :Links ok
            PFE 2 :Links ok
        FPC 5
            PFE 0 :Unused
```

In the `show chassis fabric plane` output, **FPC 5** is the AS MCC.

The term `Unused` in the output for the **show chassis fabric fpcs** and **show chassis fabric plane** command indicates that one fabric plane from each pair that share physical links (1 and 5, and 3 and 7) is inactive.

See *Junos OS System Basics and Services Command Reference* for more information.

Fabric Plane Management on MX304 Routers

IN THIS SECTION

The SFB on MX304 router supports the following functionalities:

- **Fabric Hardening:** Controls bandwidth degradation and prevents null route.

- **Fabric Fault Management:** Supported per plane. Fabric fault management per plane results in increased granularity, to identify, isolate, and repair faults.

Fabric Hardening Support and Plane Management on MX304 Routers

Fabric plane management incudes fabric hardening, that is the process to control bandwidth degradation and prevent a null route for data transmission.

MX304 routers have only one built-in SFB and line card MIC, MX304-LMIC16-BASE. The SFB has two PFEs. Each PFE supports 18 fabric planes (or sub-channels).

Table 4: LMIC support for SFB

LMICs	Switch Fabric Boards Supported	Packet Forwarding Engines (PFEs)	Fabric Planes	Fabric Redundancy
MX304-LMIC16-BASE	1 SFB	2 PFE	36	No

For details, on the fabric resiliency support, see *Fabric Plane Management on MX304 Routers*.

Table 5: Fabric Plane Management on MX304 Routers

Failure or Fault	Default Action	Configurable Action
All planes of a PFE come down (due to training failures, destination timeouts or combination of both).	Affected PFE is disabled.	Log only, FPC offline, FPC restart, FPC restart and then offline.
Multiple PFEs lose all 18 planes (number of PFEs are less than 50% in the chassis)	Affected PFEs are disabled.	Log only, FPC offline, FPC restart, FPC restart and then offline.
Combination PFEs are at fault.	Affected PFEs are disabled.	Log only, FPC offline, FPC restart, FPC restart and then offline.
All 18 planes are offlined or more than 50% of the PFEs in the chassis have faults.	SFB restart and FPC restart. If the attempt fails, PFEs are disabled.	Ignore SFB restart, Ignore FPC restart.
SFB Fatal error	SFB reset– attempts 3 times before giving up.	None

The following key CLI commands are available for fabric hardening:

- `set chassis fpc slot-number fabric bandwidth-degradation percentage`—Configures the FPC to take a specific action once bandwidth degradation reaches a certain percentage to avoid causing a null route in the chassis.
- `set chassis fabric degraded detection-enable`—Enables detection of an FPC with degraded fabric.
- `set chassis fabric degraded action-fpc-restart-disable`—Disables line card restarts to limit recovery actions from a degraded fabric condition.
- Use the commands `show chassis fabric reachability detail` to see if any fabric hardening actions are taken.
- Use command `show chassis fabric degradation` to check bandwidth information.
- Use `show chassis fabric summary extended` and `show log chassisd` for log information.

Limitations

• MX304 routers have only one built-in SFB and one FPC. Hence there is no fabric redundancy support.

• SFB offline and online is not supported. The command `request chassis sfb slot 0 {offline| online}` is not supported. You can control the operation of the specified fabric planes by using the command `request chassis fabric plane plane_number {offline| online}`.

Fabric Plane Management on MX10004 and MX10008 Devices

IN THIS SECTION

Fabric Plane Management on MX10004 and MX10008 Devices

The MX10004 router has four slots and MX10008 router has eight slots for the line cards that can support a maximum of 768 100-Gigabit Ethernet ports (4x100), 192 40-Gigabit Ethernet ports, 192 100-Gigabit Ethernet ports, or 192 400-Gigabit Ethernet ports with line card slots 0-7 that combine Packet Forwarding Engine (PFE) and Ethernet interfaces enclosed in a single assembly.

There are two models of SFBs: the JNP10004-SF or JNP10008-SF and the JNP10004-SF2 or JNP10008-SF2. SFBs installed must be of the same model type in a running chassis. On both SF and SF2 models, the SFB has eight connectors that connect to one of the eight line cards.

MX10004 and MX10008 devices support the following line cards:

- **MX10K-LC2101**—The MX10K-LC2101 line card provides a maximum bandwidth of 2.4Tbps and has six Packet Forwarding Engines (PFEs), each providing a maximum bandwidth of up to 400 Gbps.
- **MX10K-LC480**—The MX10K-LC480 line card is a Cx;7 configuration MPC with 48 ports. Each port supports a speed of 10 Gbps or 1 Gbps, providing the line card a maximum bandwidth of 480 Gbps. The MX10K-LC480 has two PFEs, each providing a maximum bandwidth of up to 240 Gbps.
- **MX10K-LC9600**—The MX10K-LC9600 is a fixed configuration 24-port line card, which provides a line rate throughput of 9.6 Tbps. The line card has twenty-four QSFP-DD ports, each capable of supporting a maximum speed of 400 Gbps.

  NOTE:
  MX10008 Hardware Guide

  The line card has 12 Packet Forwarding Engines, each providing a maximum bandwidth of 800 Gbps.
- **MX10K-LC4800**—The MX10K-LC4800 is a fixed configuration 44-port line card, which provides a line rate throughput of 4.8 Tbps. The line card has forty SFP56-56 DD ports that support 100 Gbps speed and four QSFP56-DD ports that support 400 Gbps speed.

This topic discusses the fabric plane management on these line cards.

The following Table 6 on page 70 provides information about Line Card support on SFB and SFB2.

Table 6: Line card support on SFB and SFB2

Line Cards	Switch Fabric Boards Supported	Packet Forwarding Engines (PFEs)	Fabric Planes	Fabric Redundancy
MX10K-LC2101	SFB and SFB2	6 PFE	24 (SFB), 12 (SFB2)	Yes (5+1 for SFB and SFB2)
MX10K-LC480	SFB and SFB2	2 PFE	24 (SFB), 12 (SFB2)	Yes (5+1 for SFB and SFB2)

Table 6: Line card support on SFB and SFB2 ***(Continued)***

Line Cards	Switch Fabric Boards Supported	Packet Forwarding Engines (PFEs)	Fabric Planes	Fabric Redundancy
MX10K-LC9600	SFB2	12 PFE	12	Not supported because MX10KLC9600 requires all six SFB2s to support the line rate.
MX10K-LC4800	SFB2	6 PFE	12	Not supported

NOTE: When one SFB2 fails, the line rate is achieved with 10 planes.

NOTE: MX10004 and MX10008 devices with SFB2 support interoperability of line cards.

MX10004 and MX10008 SFB2s support the following:

- Fabric fault handling: Fabric fault handling is supported per plane. Fabric fault handling per plane results in increased granularity, which helps identify, isolate, and repair faults. If an SFB has a single faulty plane, the other three planes can continue to operate. There is no need to take the entire SFB offline. For example, if a plane encounters a training failure error, the line card isolates that faulty plane; while the other planes continue to operate. Also, any cyclic redundancy check (CRC) errors on any link on the SFB are indicated on the plane, not on SFB.
- Fabric hardening: Fabric hardening is the process of controlling bandwidth degradation to prevent null route. The following key CLI commands are available for fabric hardening:
- `set chassis fpc` *`slot-number`* `fabric bandwidth-degradation percentage`—Configures the FPC to take a specific action once bandwidth degradation reaches a certain percentage to avoid causing a null route in the chassis.
- `set chassis fabric degraded detection-enable`—Enables detection of an FPC with degraded fabric.
- `set chassis fabric degraded action-fpc-restart-disable`—Disables line card restarts to limit recovery actions from a degraded fabric condition.

RELATED DOCUMENTATION

Fabric Grant Bypass | 75

Fabric Resiliency | 27

Platform Redundancy FEB Redundancy Support for High Availability of ACX7509 Devices

IN THIS SECTION

FEB redundancy is supported on the ACX7509 device. The Routing Control Board (RCB) and Forwarding Engine Board (FEB) mastership is tied and switchover together. The master Routing Engine (RE) software manages both FEBs. Graceful RE switchover (GRES) for RCB and FEB simultaneous switchover is supported with support for replication and restoration of statistics.

The graceful Routing Engine switchover (GRES) feature in Junos OS and Junos OS Evolved enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.

When GRES mode is not enabled on the Junos OS, or Junos OS Evolved, it is considered as Non-GRES mode.

Routing Engine Switchover Conditions and Prerequisites

The prerequisite for a GRES is that, the backup RCB and FEB are online for 360 seconds. The back to back switchover time is more than 360 seconds.

The switchover conditions are as follows:

- Master RCB power-failure.
- Master RE rebooted,
- Master RE plugged out,
- Master RE offlined.
- Linux kernel crash on master RE.
- Critical application failure on master RE (including PFE management, PP management & packet input output applications).
- Power fault on master FEB.

When the FEB redundancy configuration is active, the supported RCB-FEB configurations for ACX7509 are as follows:

Table 7: Supported Redundancy Modes in ACX7509

Supported Redundancy Modes	Condition
Both RCBs and FEBs present	Fully redundant system.
RCB0/FEB0 present and RCB1/FEB1 not present	Non-redundant system.
RCB1/FEB1 present and RCB0/FEB0 not present	Non-redundant system.

The unsupported redundancy modes for ACX7509 with related alarms are as follows:

Table 8: Unsupported Redundancy Modes in ACX7509

Unsupported Redundancy Modes	Condition
RCB0/FEB1 present	The unmatched FEB1 does not come online.
RCB1/FEB0 present	The unmatched FEB0 does not come online.

Table 8: Unsupported Redundancy Modes in ACX7509 *(Continued)*

Unsupported Redundancy Modes	Condition
Two RCBs with FEB0	Both RCBs and the FEB come online. RCB0/FEB0 becomes master. CLI based switchover is allowed. Fault-triggered automatic switchover does not happen.

Support for Replication and Restoration of Statistics on RE Switchover (ACX7509)

The RCB and FEB mastership is tied and switchover together. The master RE software manages both FEBs.

Traffic Flow and Switchover

The WAN traffic entering the system is forwarded to and flows through both FEBs. The FEBs route the traffic internally to the egress FPC. The FPC transmits traffic from master FEB to WAN and drops the traffic from the backup FEB. The backup FEB is in hot standby and ready to forward traffic.

Junos software supports the replication and restoration of the following statistics for traffic flow and switchover:

- Interface Statistics.
- Sub-interface Statistics (Without Queue Statistics).
- The replication for other PFE statistics entities, such as:
  - Interface per Queue Statistics
  - Firewall Statistics
  - uRPF Statistics
  - DDoS PFE Statistics (excluding the RE level DDoS counters)

Limitations for GRES mode

- RCB and FEB are a host-subsystem and switchover due to one will cause complete host sub-system switch. Hence, the independent switchover of RCB and FEB is not supported.
- For GRES, if there is sudden master PFE loss, the interface could go down briefly at the end router. The link comes back up immediately when the switchover takes effect.
- The `show interfaces` command displays the restoration of the RE based statistics.
- Ungraceful removal of FEB is not supported.

Traffic Management

When the WAN traffic is affected due to various fault scenarios, the traffic hit duration depends on:

- The time taken to detect the fault condition.
- The time taken for RCB-FEB switchover.

The RE manages the traffic switchover.

RELATED INFORMATION

Fabric Grant Bypass

IN THIS SECTION

Understanding Fabric Grant Bypass

Modular Port Concentrators (MPCs) contain one, two, or four Packet Forwarding Engines. Each Packet Forwarding Engine handles its forwarding decisions independently. Also, each Packet Forwarding Engine implements fabric queuing and flow control features required to communicate with other Packet Forwarding Engines on the same chassis. Transmitting a packet from a Packet Forwarding Engine to another involves a fabric request and grant process. As per this, the ingress Packet Forwarding Engine first sends a fabric request to the egress Packet Forwarding Engine across an active fabric plane. And when it receives the fabric grant in response, it sends the packets to the egress Packet Forwarding Engine.

However, the MX2010 and 2020 routers in some configurations are set to bypass the fabric request and grant process by default. The fabric grant bypass configuration is required to support MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFPP) on MX2020 and MX2010 platforms. On the MX Series routers with the fabric grant bypass enabled, the switch fabric takes in the fabric requests from the ingress Packet Forwarding Engine and provides fabric grants; and the ingress Packet Forwarding Engine sends the packet to the egress Packet Forwarding Engine. In this case, the switch fabric forwards the fabric request to the egress Packet Forwarding Engine, but discards the fabric grants it receives from the egress Packet Forwarding Engine.

Table 9 on page 76 describes the fabric grant bypass behavior on MX2010 and MX2020 routers.

Table 9: Fabric Grant Bypass Behavior on MX2010 and MX2020 Routers

MX Series Routers	Switch Control Board	Switch Fabric Board	Default Fabric Grant Bypass Behavior
MX2010 and MX2020	-	SFB	Enabled for all MPCs.
MX2010 and MX2020	-	SFB2	Enabled for MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFPP). Disabled for all other MPCs.

Disabling Fabric Grant Bypass to Control Congestion and Improve Performance

You can disable fabric grant bypass on the MX2020 and MX2010 routers with SFBs. Disabling the default fabric grant bypass behavior controls congestion and thus improves system behavior and performance on MX2010 and MX2020 routers. After disabling fabric grant bypass, you must reboot the router for the changes to take effect.

> **NOTE**: After you disable fabric grant bypass and reboot the router, the existing MPCs on the router where fabric grant bypass is enabled by default—such as MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16xGE-SFPP)—do not power on.

To disable fabric grant bypass to control congestion and improve system behavior and performance:

1. Disable fabric grant bypass by including the `fabric disable-grant-bypass` statement at the `[edit chassis]` hierarchy level.

   ```
   [edit chassis]
   user@host# set fabric disable-grant-bypass
   ```

2. After disabling fabric grant bypass, commit the configuration.

   ```
   [edit chassis]
   user@host# commit
   ```

   > **NOTE**: After you disable fabric grant bypass and commit the configuration, the router displays the following warning message:
   > `[edit] 'chassis' WARNING: Chassis configuration for fabric grant bypass has been changed. A system reboot is mandatory. Please reboot the system NOW. Continuing without a reboot might result in unexpected system behavior. commit complete`

3. Reboot the router for the configuration to take effect.

   ```
   user@host> request system reboot
   ```

Re-Enabling Fabric Grant Bypass

After you disable fabric grant bypass, you can re-enable it on the MX2020 and MX2010 routers with SFBs.

NOTE:

- By default, fabric grant bypass is enabled on the MX2010 and MX2020 routers.
- After you enable fabric grant bypass feature and reboot the router, the existing MPCs on the router where fabric grant bypass is enabled by default—such as MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFPP)—power on.

To re-enable fabric grant bypass:

1. Use the `delete` statement with the `fabric disable-grant-bypass` statement at the `[edit chassis]` hierarchy level to enable fabric grant bypass.

```
[edit chassis]
user@host# delete fabric disable-grant-bypass
```

2. After enabling fabric grant bypass, commit the configuration.

```
[edit chassis]
user@host# commit
```

NOTE: After you enable fabric grant bypass and commit the configuration, the router displays the following warning message:
`[edit] 'chassis' WARNING: Chassis configuration for fabric grant bypass has been changed. A system reboot is mandatory. Please reboot the system NOW. Continuing without a reboot might result in unexpected system behavior. commit complete`

3. Reboot the router for the configuration to take effect.

```
user@host> request system reboot
```

RELATED DOCUMENTATION

Smooth Upgrade from SFB to SFB2

IN THIS SECTION

Understanding the Smooth Upgrade Process

The MX2000 line of routers support Switch Fabric Board (SFB; model number: MX2000-SFB) and the enhanced Switch Fabric Board (SFB2; model number: MX2000-SFB-S). SFB2 is designed to support higher bandwidth than that provided by SFB on the MX2000 line of routers. For instance, the MX2000 line of routers with SFB support fabric bandwidth of 800 Gbps. However, the MX2000 line of routers with SFB2 can support fabric bandwidth of 1.92 Tbps. A smooth upgrade enables you to upgrade from SFB to SFB2 with minimal traffic impact on the MX2000 line of routers.

NOTE: If you have installed the Junos Continuity software package, you cannot perform a smooth upgrade from Switch Fabric Board (SFB) to Enhanced Switch Fabric Board (SFB2) on MX2010 and MX2020 routers.

This topic explains the smooth upgrade process that takes place when you upgrade from Switch Fabric Board (SFB) to enhanced Switch Fabric Board (SFB2) on MX2000 line of routers.

NOTE: The MX2000 line of routers support either SFB or SFB2 only. The MX2000 line of routers do not support SFB and SFB2 at the same time. However, during an upgrade from SFB to SFB2, the MX2000 line of routers support both SFB and SFB2 at the same time for the duration of the upgrade. But, you must upgrade all 8 SFBs to 8 SFB2s. You cannot replace 4 SFBs with 4SFB2s and retain the other SFBs.

The process of smooth upgrade from SFB and SFB2 includes the following steps:

1. Initiating the smooth upgrade process. When you initiate smooth upgrade, the router can support both SFB and SFB2 at the same time.

2. Performing the smooth upgrade. This step consists of replacing all SFBs with SFB2s.

3. Terminating the smooth upgrade. When you terminate the smooth upgrade process, the router stops supporting SFB and SFB2 at the same time.

A smooth upgrade provides the following benefits:

- The smooth upgrade eliminates network downtime during the smooth upgrade window because of 7+1 fabric redundancy. When one SFB is being upgraded to SFB2, the other seven SFBs are available to handle the traffic.

  NOTE: If multiple SFBs are upgraded at the same time, multiple fabric planes are down at any specified time and so traffic is impacted.

- When multiple fabric boards and planes come online at the same time, you can batch them together and train. This reduces the booting up time and the time taken for the plane to come online.

- On MX2000 line of routers with SFB, fabric grant bypass is enabled by default. Disabling fabric grant bypass helps control congestion and improves performance. On MX2000 line of routers with SFB, you can disable fabric grant bypass by By default, fabric grant bypass is disabled for all MPCs on MX2000 line of routers when they connect to SFB2. Fabric grant bypass is enabled by default on MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16xGE-SFPP). When fabric grant bypass is enabled by default, when those MPCs connect to SFB2, fabric grant bypass continues to be enabled and cannot be disabled.

To quickly access the information you need, click the links in Table 10 on page 80.

Table 10: Locating the Information You Need to Work on Smooth Upgrade Process

Task You Need to Perform	Where The Information Is Located
Before You begin	"Before you Begin the Smooth Upgrade Process" on page 81

Table 10: Locating the Information You Need to Work on Smooth Upgrade Process *(Continued)*

Task You Need to Perform	Where The Information Is Located
Perform a Smooth Upgrade	"Performing a Smooth Upgrade to Enhanced Switch Fabric Board (SFB2) with Minimal Impact on Traffic" on page 93

Before you Begin the Smooth Upgrade Process

Before you begin the smooth upgrade from Switch Fabric Board (SFB) to enhanced Switch Fabric Board (SFB2), complete the following tasks:

NOTE: If you have installed the Junos Continuity software package, you cannot perform a smooth upgrade from Switch Fabric Board (SFB) to Enhanced Switch Fabric Board (SFB2) on MX2010 and MX2020 routers.

- Prepare the router and install the version of Junos OS Release (16.1R1 or later) that supports the smooth upgrade process. For more information about how to install or upgrade the version of Junos OS Release, see No Link Title.
- Verify that the Switch Fabric Boards and fabric planes are online and operational. At this time, the line cards are connected to SFB.

1. To verify that all the switch fabric boards (SFBs) are online and operational, issue the following command:

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E0A50AFJ      MX2020
Midplane         REV 01   711-032387   abcd1111          Lower Backplane
Midplane 1       REV 04   711-032386   ABAB9191          Upper Backplane
PMP 1            REV 05   711-032428   ACAJ1526          Upper Power Midplane
PMP 0            REV 04   711-032426   ACAJ1585          Lower Power Midplane
FPM Board        REV 06   760-040242   ABBT8836          Front Panel Display
PSM 0            REV 01   740-050037   1EDB32101E3       DC 52V Power Supply Module
PSM 1            REV 01   740-033727   1E012130107       DC 52V Power Supply Module
PSM 2            REV 01   740-050037   1EDB3210162       DC 52V Power Supply Module
```

```
PSM 3            REV 01   740-050037   1EDB32000R6      DC 52V Power Supply Module
PSM 4            REV 01   740-050037   1EDB313005M      DC 52V Power Supply Module
PSM 5            REV 01   740-050037   1EDB321016G      DC 52V Power Supply Module
PSM 6            REV 01   740-050037   1EDB313005F      DC 52V Power Supply Module
PSM 7            REV 01   740-050037   1EDB313009X      DC 52V Power Supply Module
PSM 8            REV 01   740-050037   1EDB3130082      DC 52V Power Supply Module
PSM 9            REV 01   740-050037   1EDB32101HH      DC 52V Power Supply Module
PSM 10           REV 01   740-050037   1EDB321015G      DC 52V Power Supply Module
PSM 11           REV 01   740-050037   1EDB32101JW      DC 52V Power Supply Module
PSM 12           REV 01   740-045050   1E02224000N      DC 52V Power Supply Module
PSM 13           REV 01   740-050037   1EDB321015C      DC 52V Power Supply Module
PSM 14           REV 01   740-050037   1EDB321015J      DC 52V Power Supply Module
PSM 15           REV 01   740-045050   1E022240015      DC 52V Power Supply Module
PSM 16           REV 01   740-045050   1E02224000L      DC 52V Power Supply Module
PSM 17           REV 01   740-050037   1EDB32101EP      DC 52V Power Supply Module
PDM 1            REV 03   740-045234   1EFA3230588      DC Power Dist Module
PDM 2            REV 03   740-045234   1EFA3230508      DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009115214       RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009099720       RE-S-1800x4
CB 0             REV 23   750-040257   CAAR5968         Control Board
CB 1             REV 12   750-040257   CAAD9498         Control Board
SPMB 0           REV 02   711-041855   ABCC1066         PMB Board
SPMB 1           REV      711-041855   ABBS1488         PMB Board
SFB 0            REV 06   711-044466   ABCD4944         Switch Fabric Board
SFB 1            REV 06   711-044466   ABCD4938         Switch Fabric Board
SFB 2            REV 06   711-044466   ABCD5175         Switch Fabric Board
SFB 3            REV 06   711-044466   ABCD5160         Switch Fabric Board
SFB 4            REV 06   711-044466   ABCD4997         Switch Fabric Board
SFB 5            REV 06   711-044466   ABCD5013         Switch Fabric Board
SFB 6            REV 06   711-044466   ABCD5267         Switch Fabric Board
SFB 7            REV 06   711-044466   ABCD4968         Switch Fabric Board
FPC 0            REV 23   750-054901   CAEH6678         MPC3E NG HQoS
  CPU            REV 11   711-045719   CAEA4592         RMPC PMB
  MIC 0          REV 26   750-028392   ZM0999           3D 20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031469   17T446600017     SFP-LX10
      Xcvr 1     REV 01   740-031469   17T446600120     SFP-LX10
      Xcvr 2     REV 01   740-031469   19T446600010     SFP-LX10
      Xcvr 3     REV 01   740-031469   0ZT446600018     SFP-LX10
      Xcvr 4     REV 01   740-031469   19T446600007     SFP-LX10
      Xcvr 5     REV 01   740-031469   18T446600081     SFP-LX10
      Xcvr 6     REV 01   740-031469   18T446600088     SFP-LX10
      Xcvr 7     REV 01   740-031469   18T446600049     SFP-LX10 Xcvr 8     REV 01
```

```
740-031469   18T446600002      SFP-LX10
      Xcvr 9     REV 01   740-031469   19T446600008      SFP-LX10
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031469   18T446600032      SFP-LX10
      Xcvr 1     REV 01   740-031469   09T446600025      SFP-LX10
      Xcvr 2     REV 01   740-031469   19T446600004      SFP-LX10
      Xcvr 3     REV 01   740-031469   18T446600084      SFP-LX10
      Xcvr 4     REV 01   740-031469   18T446600060      SFP-LX10
      Xcvr 5     REV 01   740-031469   17T446600085      SFP-LX10
      Xcvr 6     REV 01   740-031469   17T446600014      SFP-LX10
      Xcvr 7     REV 01   740-031469   17T446600315      SFP-LX10
      Xcvr 8     REV 01   740-031469   18T446600043      SFP-LX10
      Xcvr 9     REV 01   740-031469   0ZT446600017      SFP-LX10
  MIC 1          REV 19   750-033199   CAAJ1818          1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN           1X100GE CFP
FPC 1            REV 32   750-028467   ZR1986            MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ZT7025            AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMH0285           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 1       REV 01   740-031980   AHK011H           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   APK0569           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
FPC 2            REV 04   750-044444   ZA7865            MPCE Type 2 3D P
  CPU            REV 02   711-038484   ZB2728            MPCE PMB 2G
  MIC 0          REV 07   750-028390   XY2158            3D 40x 1GE(LAN) RJ45
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) RJ45
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) RJ45
    PIC 2                 BUILTIN      BUILTIN           10x 1GE(LAN) RJ45
    PIC 3                 BUILTIN      BUILTIN           10x 1GE(LAN) RJ45
  MIC 1
  QXM 0          REV 05   711-028408   ZC3420            MPC QXM
  QXM 1          REV 05   711-028408   ZC3350            MPC QXM
FPC 3            REV 22   750-054564   CADG6972          MPC5E 3D 2CGE+4XGE
  CPU            REV 11   711-045719   CADC7599          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   193363A00483      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   1YT517101829      SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0       REV 01   740-052505   XUF0GPX           CFP2-100G-SR10
  PIC 2                   BUILTIN      BUILTIN           2X10GE SFPP OTN
  PIC 3                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
```

```
FPC 6           REV 11   750-045372   CABT0840         MPCE Type 3 3D
  CPU           REV 08   711-035209   CABL0889         HMPC PMB 2G
MIC 0         REV 27   750-028392   CABR4723         3D 20x 1GE(LAN) SFP
    PIC 0                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0    REV 01   740-011782   P9229UM          SFP-SX
      Xcvr 1    REV 01   740-011782   P9P0X6V          SFP-SX
      Xcvr 2    REV 01   740-011613   PCE01W5          SFP-SX
      Xcvr 4    REV 01   740-011613   PD63DEN          SFP-SX
      Xcvr 5    REV 02   740-011613   PG12FSF          SFP-SX
      Xcvr 7    REV 01   740-011782   PCL3UDY          SFP-SX
      Xcvr 8    REV 01   740-011613   PE713Z9          SFP-SX
      Xcvr 9    REV 01   740-011613   AM0846SAQA5      SFP-SX
    PIC 1                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0    REV 01   740-011613   P9F16KE          SFP-SX
      Xcvr 1    REV 01   740-031851   AM1045SU91U      SFP-SX
      Xcvr 4    REV 01   740-011613   PAJ4SY8          SFP-SX
      Xcvr 5    REV 01   740-011782   P9228K7          SFP-SX
  MIC 1         REV 27   750-028392   CABT5724         3D 20x 1GE(LAN) SFP
    PIC 2                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0    REV 02   740-011613   AM0925SBG5T      SFP-SX
      Xcvr 1             NON-JNPR     P7K1PUX          SFP-SX
      Xcvr 2    REV 01   740-011613   PFF2DHH          SFP-SX
      Xcvr 4    REV 01   740-011613   PD63DF2          SFP-SX
      Xcvr 5    REV 02   740-011613   AM1033SH3DH      SFP-SX
      Xcvr 6    REV 01   740-011613   PE70W8W          SFP-SX
      Xcvr 9    REV 01   740-011613   PD62W9W          SFP-SX
    PIC 3                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0    REV 02   740-013111   9154876          SFP-T
      Xcvr 2    REV 01   740-011613   AM0846SAQ9H      SFP-SX
      Xcvr 5    REV 01   740-011613   AM0820S9T2C      SFP-SX
      Xcvr 9    REV 01   740-011613   AM0805S8LGQ      SFP-SX
FPC 7           REV 27   750-033205   ZL6014           MPCE Type 3 3D
  CPU           REV 07   711-035209   ZK9068           HMPC PMB 2G
  MIC 0         REV 04   750-028392   JR6231           3D 20x 1GE(LAN) SFP
    PIC 0                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0    REV 01   740-031851   AM1045SU93A      SFP-SX
    PIC 1                BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 4    REV 01   740-011782   P9P1050          SFP-SX
      Xcvr 9    REV 01   740-011613   PFF2K74          SFP-SX
  MIC 1         REV 19   750-033199   CAAF0016         1X100GE CFP
    PIC 2                BUILTIN      BUILTIN          1X100GE CFP
FPC 11          REV 16   750-037358   CAAL1014         MPC4E 3D 32XGE
  CPU           REV 08   711-035209   CAAS2637         HMPC PMB 2G
```

```
  PIC 0
  PIC 1
  PIC 2
  PIC 3                            FPC 12          REV 29   750-031090
ZA1887           MPC Type 2 3D EQ
  CPU            REV 06   711-030884   YR6876          MPC PMB 2G
FPC 13           REV 36   750-056519   CAFW4205        MPC7E 3D MRATE-12xQSFPP-XGE-XLGE-CGE
  CPU            REV 16   750-057177   CAFY5688        SMPC PMB
  PIC 0                   BUILTIN      BUILTIN         MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-054053   QF3208FT        QSFP+-4X10G-SR
    Xcvr 3       REV 01   740-032986   QB171000        QSFP+-40G-SR4
    Xcvr 5       REV 01   740-058732   1CJQA10700C     QSFP-100GBASE-LR4
  PIC 1                   BUILTIN      BUILTIN         MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-054053   QF3208G2        QSFP+-4X10G-SR
    Xcvr 1       REV 01   740-054053   QF3208G3        QSFP+-4X10G-SR
    Xcvr 2                NON-JNPR     F2M2010439      QSFP-100GBASE-LR4
    Xcvr 3       REV 01   740-046565   QF3300ZQ        QSFP+-40G-SR4
    Xcvr 5       REV 01   740-058734   1ACQ104202U     QSFP-100GBASE-SR4
FPC 14           REV 68   750-044130   ABDC2916        MPC6E 3D
  CPU            REV 12   711-045719   ABDC2710        RMPC PMB
FPC 16           REV 22   750-037355   CABW1289        MPC4E 3D 2CGE+8XGE
  CPU            REV 08   711-035209   CABR9796        HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN         4x10GE SFPP
  PIC 1                   BUILTIN      BUILTIN         1X100GE CFP
  PIC 2                   BUILTIN      BUILTIN         4x10GE SFPP
  PIC 3                   BUILTIN      BUILTIN         1X100GE CFP
FPC 17           REV 23   750-037355   CACL2280        MPC4E 3D 2CGE+8XGE
  CPU            REV 10   711-035209   CACK9073        HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN         4x10GE SFPP
  PIC 1                   BUILTIN      BUILTIN         1X100GE CFP
  PIC 2                   BUILTIN      BUILTIN         4x10GE SFPP
  PIC 3                   BUILTIN      BUILTIN         1X100GE CFP
FPC 18           REV 23   750-054901   CAEV3700        MPC3E NG HQoS
  CPU            REV 12   711-045719   CAFK4017        RMPC PMB
  MIC 0          REV 19   750-033199   CAAJ9717        1X100GE CFP
    PIC 0                 BUILTIN      BUILTIN         1X100GE CFP
  MIC 1          REV 15   750-033199   ZP6432          1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN         1X100GE CFP
FPC 19           REV 29   750-063414   CAEJ2194        MPC9E 3D
  CPU            REV 02   750-057177   CACN2561        SMPC PMB
  MIC 0          REV 01   750-055992   CADV4595        MRATE-12xQSFPP-XGE-XLGE-CGE
    PIC 0                 BUILTIN      BUILTIN         MRATE-12xQSFPP-XGE-XLGE-CGE
      Xcvr 0     REV 01   740-046565   QF3300ZG        QSFP+-40G-SR4
```

```
      Xcvr 1     REV 01   740-046565   QF330122         QSFP+-40G-SR4
      Xcvr 2     REV 01   740-046565   QF33011P         QSFP+-40G-SR4
      Xcvr 3     REV 01   740-046565   QF3300ZU         QSFP+-40G-SR4
      Xcvr 4     REV 01   740-046565   QF3300ZS         QSFP+-40G-SR4
      Xcvr 5     REV 01   740-046565   QF3300ZN         QSFP+-40G-SR4
      Xcvr 6     REV 01   740-046565   QF3300ZP         QSFP+-40G-SR4
            Xcvr 7     REV 01   740-046565   QF3300ZT         QSFP+-40G-SR4
      Xcvr 8     REV 01   740-046565   QF3300ZM         QSFP+-40G-SR4
      Xcvr 9     REV 01   740-046565   QF3300ZR         QSFP+-40G-SR4
      Xcvr 10    REV 01   740-046565   QF330105         QSFP+-40G-SR4
      Xcvr 11    REV 01   740-046565   QF3300ZK         QSFP+-40G-SR4
  MIC 1          REV 08   750-055992   CAEX1421         MRATE-12xQSFPP-XGE-XLGE-CGE
    PIC 1                 BUILTIN      BUILTIN          MRATE-12xQSFPP-XGE-XLGE-CGE
      Xcvr 6     REV 01   740-046565   QF330100         QSFP+-40G-SR4
ADC 0            REV 19   750-043596   ABCK6658         Adapter Card
ADC 1            REV 17   750-043596   ABCB7201         Adapter Card
ADC 2            REV 05   750-043596   CAAC2076         Adapter Card
ADC 3            REV 13   750-043596   ABBX5549         Adapter Card
ADC 6            REV 17   750-043596   ABCB7226         Adapter Card
ADC 7            REV 01   750-043596   ZV4079           Adapter Card
ADC 11           REV 17   750-043596   ABCD5472         Adapter Card
ADC 12           REV 17   750-043596   ABCB7147         Adapter Card
ADC 13           REV 17   750-043596   ABCD5410         Adapter Card
ADC 16           REV 17   750-043596   ABCB7047         Adapter Card
ADC 17           REV 17   750-043596   ABCD5525         Adapter Card
ADC 18           REV 17   750-043596   ABCD5391         Adapter Card
Fan Tray 0       REV 01   760-042349   ACAY4801         FanTray v2
Fan Tray 1       REV 01   760-042349   ACAY4802         FanTray v2
Fan Tray 2       REV 01   760-042349   ACAY4803         FanTray v2
Fan Tray 3       REV 01   760-042349   ACAY4800         FanTray v2
```

2. To verify that all the fabric planes are available and operational, issue the following command:

```
user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
```

```
            PFE 2 :Links ok
            PFE 3 :Links ok
        FPC 2
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 3
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 6
            PFE 0 :Links ok
        FPC 7
            PFE 0 :Links ok
        FPC 11
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 12
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 13
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 14
            PFE 0 :Links ok
            PFE 1 :Links ok
            PFE 2 :Links ok
            PFE 3 :Links ok
        FPC 16
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 17
            PFE 0 :Links ok
            PFE 1 :Links ok
        FPC 18
            PFE 0 :Links ok
        FPC 19
            PFE 0 :Links ok
            PFE 1 :Links ok
     PFE 2 :Links ok
            PFE 3 :Links ok
  Plane 1
    Plane state: ACTIVE
        FPC 0
            PFE 0 :Links ok
```

```
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 3
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 6
        PFE 0 :Links ok
    FPC 7
        PFE 0 :Links ok
    FPC 11
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 12
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 13
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 14
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 16
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 17
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 18
        PFE 0 :Links ok
    FPC 19
        PFE 0 :Links ok
    PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
...
```

```
Plane 7
  Plane state: ACTIVE
      FPC 0
          PFE 0 :Links ok
      FPC 1
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
      FPC 2
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 3
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 6
          PFE 0 :Links ok
      FPC 7
          PFE 0 :Links ok
      FPC 11
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 12
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 13
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 14
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
      FPC 16
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 17
          PFE 0 :Links ok
          PFE 1 :Links ok
      FPC 18
          PFE 0 :Links ok
      FPC 19
```

```
          PFE 0 :Links ok
          PFE 1 :Links ok
          PFE 2 :Links ok
          PFE 3 :Links ok
```

3. To verify that the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the Switch Fabric Boards (SFBs) are eligible for carrying traffic, issue the following command:

```
user@host>show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 1
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #1
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #2
      Plane 0: Plane enabled
```

```
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #3
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
            Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
  FPC 2
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
  ...
  FPC 19
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
```

```
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #2
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #3
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
```

4. To verify the state of all fabric planes and the elapsed time, issue the following command:

```
user@host> show chassis fabric summary
Plane   State    Uptime
 0      Online   11 hours, 13 minutes, 27 seconds
 1      Online   11 hours, 13 minutes, 6 seconds
 2      Online   11 hours, 12 minutes, 45 seconds
 3      Online   11 hours, 12 minutes, 24 seconds
 4      Online   11 hours, 12 minutes, 2 seconds
 5      Online   11 hours, 11 minutes, 41 seconds
 6      Online   11 hours, 11 minutes, 20 seconds
```

```
7      Online   11 hours, 10 minutes, 59 seconds

Note: For extended summary, use
      show chassis fabric summary extended
```

Performing a Smooth Upgrade to Enhanced Switch Fabric Board (SFB2) with Minimal Impact on Traffic

IN THIS SECTION

This example shows how to perform a smooth upgrade from the Switch Fabric Board (SFB) to the enhanced Switch Fabric Board (SFB2) on the MX2000 line of routers. A smooth upgrade helps reduce network downtime because of 7+1 fabric redundancy. When one SFB is being upgraded to SFB2, the other 7 SFBs are available to handle the traffic.

NOTE: On MX2010 and MX2020 routers, if you have installed the Junos Continuity software package or if the router is not configured to allow multiple versions of the SFBs to coexist, you cannot perform a smooth upgrade from SFB to SFB2.

When not using a smooth upgrade, use one of the following methods to upgrade to SFB2:

- Power off the router, replace the SFB with SFB2, and then power on the router.
- Take both the Routing Engines offline, replace the SFB with SFB2, and then bring both the Routing Engines online.

Requirements

This example uses the following hardware and software components:

- MX2020 router with dual Routing Engines

- 8 Switch Fabric Boards (SFBs)
- 8 enhanced Switch Fabric Boards (SFB2s)
- Junos OS Release 16.1R1 or later release

Before you begin the smooth upgrade, ensure that you:

- Prepare the router and install the version of Junos OS Release that supports the enhanced Switch Fabric Board (SFB2).
- Verify that the existing SFBs are online and operational and also check the status of the fabric planes.

For more information about what you must do before you commence smooth upgrade, see "Before you Begin the Smooth Upgrade Process" on page 81.

Overview

IN THIS SECTION

The smooth upgrade process is used to upgrade from Switch Fabric Board (SFB) to enhanced Switch Fabric Board (SFB2) with minimal traffic impact. The existing SFBs are replaced one by one, in any order, by the new SFB2s. Because you are replacing a single SFB at a time, the remaining SFBs handle the traffic and so there is minimal impact to traffic. SFB2 is supported only on MX2020 and MX2010 routers.

Topology

This example shows how to perform a smooth upgrade on an MX2020 router that has eight SFBs. The 8 SFBs are replaced with 8 enhanced switch fabric boards (SFB2). First, initiate the smooth upgrade process and then take a single SFB offline. Replace the SFB with an SFB2, and then bring the SFB2 online. You can then repeat the steps for the other seven SFBs.

After you upgrade all the SFBs to SFB2s, the fabric bandwidth per slot of MPC8E and MPC9E on the MX2020 router is increased from 11 Gbps to 25 Gbps. However, the upgrade does not impact the fabric bandwidth per slot of MPC7.

Configuration

IN THIS SECTION

To upgrade from SFB to SFB2, perform the following tasks:

Initiating the Smooth Upgrade Process

Step-by-Step Procedure

By default, the MX2000 line of routers do not support both SFB and SFB2 at the same time. However, when you initiate the smooth upgrade process, the router can support both SFB and SFB2 at the same time. So, before you replace an SFB with an SFB2, you must initiate the smooth upgrade process.

1. In configuration mode, at the `[edit]` hierarchy level, Initiate the smooth upgrade process for the SFBs.

```
[edit]
user@host# set chassis state sfb-upgrade on
```

2. Commit the changes by using the `commit` statement and exit the configuration mode.

```
[edit]
user@host# commit
```

3. In operational mode, verify that you have initiated the smooth upgrade process.

```
user@host> show configuration chassis
state {
sfb-upgrade on;
}
```

Performing the Smooth Upgrade

Step-by-Step Procedure

1. In operational mode, take the SFBs offline. There is no specific order that needs to be maintained. In this example, you start with the SFB in slot 7 first.

```
user@host> request chassis sfb slot 7 offline
```

2. Verify that the SFB is offline.

```
user@host> show chassis sfb
```

```
Slot  State            Uptime
 0    Online            1 day, 12 hours, 6 minutes, 59 seconds
 1    Online            1 day, 12 hours, 6 minutes, 37 seconds
 2    Online            1 day, 12 hours, 6 minutes, 16 seconds
 3    Online            1 day, 12 hours, 5 minutes, 55 seconds
 4    Online            1 day, 12 hours, 5 minutes, 33 seconds
 5    Online            1 day, 12 hours, 5 minutes, 12 seconds
 6    Online            1 day, 12 hours, 4 minutes, 51 seconds
 7    Offline            --- Offlined by cli command ---
```

3. Replace the SFB that is offline with the enhanced SFB (SFB2). Minimal traffic loss is expected as only a single SFB is replaced and other seven SFBs are operational and handle the traffic.

4. In operational mode, bring the SFB2 online.

```
user@host> request chassis sfb slot 7 online
```

5. Verify that the SFB2 is online.

```
user@host> show chassis sfb
```

```
Slot  State            Uptime
 0    Online            1 day, 12 hours, 16 minutes, 38 seconds
 1    Online            1 day, 12 hours, 16 minutes, 16 seconds
```

```
2     Online             1 day, 12 hours, 15 minutes, 55 seconds
3     Online             1 day, 12 hours, 15 minutes, 34 seconds
4     Online             1 day, 12 hours, 15 minutes, 12 seconds
5     Online             1 day, 12 hours, 14 minutes, 51 seconds
6     Online             1 day, 12 hours, 14 minutes, 30 seconds
7     Online             38 seconds
```

6. Repeat Step 3 through Step 5 for upgrading the other SFBs. We recommend that you upgrade fabric boards one at a time for minimal traffic impact.

Terminating the Smooth Upgrade Process

Step-by-Step Procedure

After all the SFBs are upgraded to the enhanced SFB (SFB2), you can terminate the smooth upgrade process. When the smooth upgrade process is initiated, SFB and SFB2 can coexist on the same router. When you terminate the smooth upgrade process, the router can have only SFB or SFB2 and not both at the same time.

1. In configuration mode, at the `[edit]` hierarchy level, terminate the smooth upgrade process.

   NOTE: You can also use the `delete chassis state sfb-upgrade` command to terminate the smooth upgrade process.

```
[edit]
user@host# set chassis state sfb-upgrade off
```

2. Commit the changes by using the `commit` statement and exit configuration mode.

```
[edit]
user@host# commit
```

3. In operational mode, verify that you have initiated the smooth upgrade process.

```
user@host> show configuration chassis
state {
sfb-upgrade off;
}
```

Verification

IN THIS SECTION

To confirm that you have upgraded SFB to SFB2 on the MX2020 router, perform these tasks:

Verifying That the Switch Fabric Board (SFB) is Offline

Purpose

To verify that the SFB on a particular slot, for instance slot 1, is offline.

Action

From operational mode, enter the `show chassis fabric fpcs` command.

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 2
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane disabled >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 4
  PFE #0
          Plane 0: Plane enabled
      Plane 1: Plane disabled >>>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
```

```
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane disabled  >>>>
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane disabled  >>>>
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane disabled  >>>>>
    Plane 2: Plane enabled
        Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 6
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane disabled   >>>>>
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane disabled   >>>>>
```

```
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
  FPC 7
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane disabled  >>>>>
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane disabled  >>>>>
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #2
        Plane 0: Plane enabled
        Plane 1: Plane disabled  >>>>>
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #3
        Plane 0: Plane enabled
        Plane 1: Plane disabled >>>>>>
        Plane 2: Plane enabled
        Plane 3: Plane enabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
```

```
      Plane 7: Plane enabled
```

From operational mode, enter the `show chassis fabric summary` command.

```
user@host> show chassis fabric summary
Plane   State    Uptime
 0      Online   3 minutes, 14 seconds
 1      Offline
 2      Online   1 hour, 56 minutes, 53 seconds
 3      Online   1 hour, 56 minutes, 39 seconds
 4      Online   1 hour, 56 minutes, 25 seconds
 5      Online   1 hour, 56 minutes, 11 seconds
 6      Online   1 hour, 55 minutes, 56 seconds
 7      Online   1 hour, 42 minutes, 28 seconds

 Note: For extended summary, use
       show chassis fabric summary extended
```

Meaning

The SFB in Slot 1 has been taken offline.

Verifying That the Enhanced Switch Fabric Board (SFB2) is Online

Purpose

To verify that the enhanced switch fabric board (SFB2) inserted in the same slot (slot 1) is online.

Action

From operational mode, enter the `show chassis fabric fpcs` command.

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 2
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled >>>>>>
      Plane 2: Plane enabled
```

```
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 4
  PFE #0
          Plane 0: Plane enabled
      Plane 1: Plane enabled >>>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #1
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #2
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #3
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>>
      Plane 2: Plane enabled
          Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 6
```

```
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled   >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #1
      Plane 0: Plane enabled
      Plane 1: Plane enabled   >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
FPC 7
  PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #1
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #2
      Plane 0: Plane enabled
      Plane 1: Plane enabled  >>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
```

```
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
  PFE #3
      Plane 0: Plane enabled
      Plane 1: Plane enabled >>>>>>
      Plane 2: Plane enabled
      Plane 3: Plane enabled
      Plane 4: Plane enabled
      Plane 5: Plane enabled
      Plane 6: Plane enabled
      Plane 7: Plane enabled
```

From operational mode, enter the `show chassis fabric summary` command.

```
user@host> show chassis fabric summary
Plane   State    Uptime
 0      Online   6 minutes, 38 seconds
 1      Online   2 minutes, 12 seconds >>>>
 2      Online   2 hours, 17 seconds
 3      Online   2 hours, 3 seconds
 4      Online   1 hour, 59 minutes, 49 seconds
 5      Online   1 hour, 59 minutes, 35 seconds
 6      Online   1 hour, 59 minutes, 20 seconds
 7      Online   1 hour, 45 minutes, 52 seconds
```

Meaning

The SFB2 that replaced the SFB on slot 1 is online and operational.

RELATED DOCUMENTATION

Understanding Fabric Grant Bypass | 76

Disabling Fabric Grant Bypass to Control Congestion and Improve Performance | 77

4

CHAPTER

Power Management

Configuring Ambient Temperature

IN THIS SECTION

Understanding How Configuring Ambient Temperature Helps Optimize Power Utilization

The key to managing power in network infrastructure is the efficient utilization of provisioned power. Provisioned power is the minimum power that is required to bring a router or a switch online. Junos OS determines the minimum required power by considering the worst-case power requirement for all the FRUs installed in the device. One of the methods to optimize the provisioned power on MX Series routers or EX9200 switches is to configure the device to operate at a cooler temperature. You can enable a device to operate at a lower operating temperature by configuring a lower ambient temperature.

Ambient temperature is the maximum operating temperature for a device. By configuring an ambient temperature, you can optimize power provisioned for the cooling system and the line cards. The maximum speed at which fans operate depends on the configured ambient temperature. As the fan speed increases, the power consumed by the fans increases. As a result, the device consumes more power when the temperature is high because the fans run faster to maintain the operating temperature of the chassis within the configured limits.

When a router or a switch restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If enough power is not available, a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature. You can configure a new higher ambient temperature only after you make more power available by adding new power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

A specific ambient temperature value might not be applicable to a different geographical location, for example, in a colder region. For devices operating in colder regions, you can configure a lower ambient temperature, which helps reduce provisioned power significantly. However, in a region of higher temperature, you might need to configure a higher ambient temperature to ensure smooth functioning of the device. For example, if the router or switch operates in a colder region, you can set the ambient temperature to 25°C, which reduces the maximum fan speed, thereby reducing the maximum power consumption. Thus, by configuring an appropriate ambient temperature, you can reduce the provisioned power and save cost on network power infrastructure.

You can configure ambient temperature by using the `set chassis ambient-temperature (25C|40C|55C)` statement at the `[edit chassis]` hierarchy level. The default ambient temperature for MX Series routers and EX9200 switches is 40°C.

> **NOTE**: The `set chassis ambient-temperature` *`temperature`* and `show chassis ambient-temperature` commands are not supported on PTX10004, PTX10008, and PTX10016 routers because their power allocation is based on the default ambient temperature.

Monitoring the Power Consumption of PTX5000 FPCs by Configuring the Ambient Temperature

You can configure the ambient temperature of the PTX5000 chassis to manage power allocated to the FPCs. You can set the ambient temperature of the chassis at 25° C, or 40° C. On system initialization, the power manager reads the ambient temperature and allocates power to the FPCs according to the power budget policy at that temperature.

1. To configure the ambient temperature, include the **set chassis ambient-temperature 25|40|55** statement at the `[edit]` hierarchy leve in the configuration mode:

   ```
   [edit]
       user@host# set chassis ambient-temperature 25|40
   ```

2. To verify the ambient temperature of the chassis, use the **show chassis ambient-temperature** command at the `[edit]` hierarchy level in the operational mode:

   ```
   [edit]
       user@host> show chassis ambient-temperature
   ```

   ```
   Ambient Temperature: 25C
   ```

To verify the power consumption of the FPCs, use the following statements:

1. Use the **show chassis power detail | grep "FPC"** statement at the `[edit]` hierarchy level to view the power consumption of the FPCs.

```
user@host> show chassis power detail |
grep "FPC"

     FPC 0                 448
  FPC 1                 419
  FPC 2                 373
  FPC 3                   0
  FPC 4                   0
  FPC 5                   0
  FPC 6                   0
  FPC 7                   0
```

   Alternatively use the SNMP MIB command, `show snmp mib walk jnxOperatingFRUPower | grep "\.7\."` to view the power consumption of each FPC:

```
user@host> show snmp mib walk jnxOperatingFRUPower
| grep "\.7\."

jnxOperatingFRUPower.4.1.7.0 = 0
jnxOperatingFRUPower.7.1.0.0 = 457              < ------  For FPC 0
jnxOperatingFRUPower.7.2.0.0 = 428              < ------  For FPC 1
jnxOperatingFRUPower.7.3.0.0 = 381              < ------  For FPC 2
jnxOperatingFRUPower.15.7.0.0 = 0
```

2. Use the **show chassis alarms** statement to view the alarms generated for any of the FPCs:

```
user@host> show chassis alarms

Alarm time               Class  Description
2007-04-08 05:51:12 UTC  Minor  FPC 1, Consumption > 90percent of allocated Budget
2007-04-08 05:51:12 UTC  Minor  FPC 0, Consumption > 90percent of allocated Budget
2007-04-08 05:50:26 UTC  Minor  FPC 0 SIB Link Error
2007-04-08 05:49:34 UTC  Minor  SIB 0 FPC Link Error
2007-04-08 05:48:02 UTC  Minor  No Redundant Power for FPC 0-7
2007-04-08 05:48:01 UTC  Minor  No Redundant Power for Rear Chassis
2007-04-08 05:48:01 UTC  Minor  No Redundant Power for Fan 0-2
```

If an FPC consumes more than 90% of the allocated power budget, the `Consumption > 90percent of allocated Budget` alarm is raised. FPC power consumption is measured at intervals of 65 seconds.

NOTE: Starting in Junos OS Release 18.4R1, the PTX5000 routers do not raise a chassis alarm in the following events:

- Power consumption by an FPC exceeds 90% of the allocated power budget.
- Power consumption by an FPC exceeds 100% of the allocated power budget (in this case, a system log is registered).

NOTE: If the PTX5000 chassis has redundant power supply modules, and if one PSM fails, the FPCs can still be online. Only the `No redundant power supply` alarm is raised.

If the PTX5000 chassis does not have redundant power supply modules, failure of one PSM can cause the FPCs to go offline, depending on the total chassis power available at that time.

3. When the power consumption of an FPC is more than the allocated budget for three consecutive intervals, the `Consumption > 90percent of allocated Budget` is cleared and `PWR Range Overshoot` alarms is raised for that particular FPC and the ambient temperature is set to the next higher setting.

```
user@host> show chassis alarms

9 alarms currently active
Alarm time               Class  Description
2007-04-08 05:56:38 UTC  Minor  FPC 2, Consumption > 90percent of allocated Budget
2007-04-08 05:55:33 UTC  Minor  FPC 1, PWR Range Overshoot
2007-04-08 05:53:22 UTC  Minor  FPC 0, PWR Range Overshoot
2007-04-08 05:50:26 UTC  Minor  FPC 0 SIB Link Error
2007-04-08 05:49:34 UTC  Minor  SIB 0 FPC Link Error
2007-04-08 05:48:02 UTC  Minor  No Redundant Power for FPC 0-7
2007-04-08 05:48:01 UTC  Minor  No Redundant Power for Rear Chassis
2007-04-08 05:48:01 UTC  Minor  No Redundant Power for Fan 0-2
```

NOTE: `Consumption > 90percent of allocated Budget` alarms are updated according to the new ambient temperature setting but the chassis ambient temperature is not changed.

```
user@host> show chassis alarms
5 alarms currently active
Alarm time              Class  Description
2007-04-01 04:36:53 UTC  Minor  No Redundant Power for FPC 0-7
2007-04-01 04:36:52 UTC  Minor  No Redundant Power for Rear Chassis
2007-04-01 04:36:51 UTC  Minor  No Redundant Power for Fan 0-2
2007-04-01 04:36:47 UTC  Minor  PDU 1 Absent
```

a. You can verify the temperature by using the **show chassis ambient-temperature** command.

```
user@host> show chassis ambient-temperature
Ambient Temperature: 25C
```

b. Enter the configuration mode and check the configured ambient temperature. Use the **show chassis ambient temperature** operational mode command.

```
user@host# show chassis ambient temperature
Ambient Temperature: 25C
```

This is set to the last configured value.

c. To clear the temperature set for the overshooting condition, use the **request chassis power-manager reset ambient-config** command.

```
user@host> request chassis power-manager
reset ambient-config
```

Verify the ambient temperature after the reset.

```
show chassis ambient-temperature
Ambient Temperature: 25C
```

4. Verify the active alarms in the chassis by using the `show chassis alarms` command.

```
user@host> show chassis alarms
7 alarms currently active
Alarm time               Class  Description
2007-04-01 04:36:53 UTC  Minor  No Redundant Power for FPC 0-7
2007-04-01 04:36:52 UTC  Minor  No Redundant Power for Rear Chassis
2007-04-01 04:36:51 UTC  Minor  No Redundant Power for Fan 0-2
2007-04-01 04:36:47 UTC  Minor  PDU 1 Absent
```

Managing Power Allocated to PTX5000 FPCs on the Basis of Chassis Ambient Temperature Configuration

The power management feature of the PTX5000 Packet Transport Router is enhanced to manage the power supplied to the FPCs on the router by configuring the ambient temperature of the chassis. You can set the ambient temperature of the chassis at 25° C, or 40° C. On system initialization, the power manager reads the ambient temperature and allocates power to the FPC according to the power budget policy at that temperature. If the actual power consumption of any FPC exceeds the configured value for more than three minutes, the power manager overrides the configured ambient temperature setting of that FPC, and resets its ambient temperature to the next higher level and reallocates power according to the new temperature setting. All the overshooting FPCs remain in the dynamic ambient temperature mode until the next reboot, or until you override it with a CLI command. The power manager then resets the power budget of the FRUs according to the configured ambient temperature setting.

NOTE: If the ambient temperature is not set, then, 55° C is considered as the default ambient-temperature and FPCs are assigned power according to the default ambient temperature.

For example, if the chassis ambient temperature is set to 25° C, the power manager allocates power to the FPCs according to the power budget policy at 25 ° C. If an FPC consumes more than 90% of the allocated power, an alarm—`Consumption > 90percent of allocated Budget`—is raised. If the FPC power consumption exceeds the allocated power for more than three minutes, the `PWR Range Overshoot` alarm is raised and the power manager reallocates power to that FPC according to the next higher temperature setting, that is, 40° C .

NOTE: During the `PWR Range Overshoot` alarm condition, you cannot reconfigure or delete the ambient temperature setting. You can reset the ambient temperature to the earlier setting after clearing the alarm condition by using the **request chassis power-manager reset ambient-config** command.

NOTE: If the PTX5000 chassis has redundant power supply modules, and if one PSM fails, the FPCs can still be online. Only the `No redundant power supply` alarm is raised.

If the PTX5000 chassis does not have redundant power supply modules, failure of one PSM can cause the FPCs to go offline, depending on the total chassis power available at that time.

RELATED DOCUMENTATION

ambient-temperature

show chassis temperature-thresholds

Understanding How Dynamic Power Management Enables Better Utilization of Power | **113**

Managing Power

IN THIS SECTION

Understanding How Dynamic Power Management Enables Better Utilization of Power

You can use the dynamic power management feature to better utilize the power available in the power entry module (PEM). Whether or not a new hardware component is powered on depends on the availability of power in the PEM. A component is not powered on if the PEM cannot meet the worst-case power requirement for that component. Starting in Junos OS Release 15.1R1, MX Series routers support dynamic power management. Starting in Junos OS Release 17.2R1, EX9200 switches support dynamic power management.

The maximum power that each type of MIC consumes is maintained in a static database. The chassis daemon process (`chassisd`), which manages power budgeting for all line cards, uses this data when budgeting power for MICs. MICs are brought online only after the chassis daemon verifies that the worst-case power required for the MICs and the power required for all the online FRUs (Field Replaceable Units: Replaceable or swappable Junos device and device parts) are available in the PEM.

In Junos OS Release 15.1R1, for MX Series routers, dynamic power management for MICs is disabled by default. You can enable the feature by enabling the `mic-aware-power-management` statement at the `[edit chassis]` hierarchy level. When dynamic power management is disabled, the chassis daemon checks for the worst-case power requirement of the MPC and the MICs before allocating power for the MPC. Whereas, when `mic-aware-power-management` statement is enabled, the chassis daemon considers the power requirement of only the MPCs. The worst-case power consumption by the MICs is not considered while the chassis daemon budgets power for the MPC. Power budgeting for MICs is done only after the MPC is powered on and the MICs come online. Every time you disable or enable dynamic power management, you must restart the chassis or the MPC for the changes to take effect.

In Junos OS Release 17.2R1, for EX9200 switches, dynamic power management for MICs is enabled by default.

Starting from Junos OS Release 17.3R1, for MX10003 routers, mic-aware dynamic power management is enabled by default.

Starting from Junos OS Release 18.2R1, for JNP10K-LC2101 MPC on MX10008 routers, dynamic power management is enabled by default. However, dynamic power management for MICs is not supported on JNP10K-LC2101 because JNP10K-LC2101 is a fixed configuration MPC and supports only built-in PICs.

After you enable the dynamic power management feature, use the `set chassis preserve-fpc-poweron-sequence` configuration mode command to preserve the sequence in which MPCs are powered on. This configuration is required to maintain the order in which the MPCs come online after a router or switch restart.

NOTE: In Junos OS Release 15.1F5 and later, dynamic power management is enabled by default on several MPCs. Models include MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, MPC2E-3D-NG-Q, MPC6E, MPC7E-MRATE, and MPC7E-10G on MX240, MX480, MX960, MX2010, and MX2020 and on MPC8E and MPC9E on MX2010, and MX2020 Universal Routing Platforms.

Understanding Power Management on the PTX5000

IN THIS SECTION

Starting in Junos OS Release 14.1, the power management feature for PTX5000 routers ensures that at any time, the chassis power requirements do not exceed the available chassis power. The PTX5000 has two PDUs to meet the power requirements of the chassis. Each PDU is capable of providing power to the chassis on its own. In case the power requirement exceeds the individual capacity of a PDU, the required power is provided by both the PDUs and the `No redundant power supply` alarm is triggered. If the system cannot provide power for all the installed FPCs or PICs, the system brings down FPCs or PICs that in can no longer provide power for and the `Insufficient Power - FRU(s) went offline` alarm is raised.

The power management feature provides the following functionality:

- Power management ensures that high-priority FPCs continue to receive power when the system does not have sufficient power to keep all the FPCs online.

- Power management ensures that if a power supply fails, the router can continue to operate normally by keeping high-priority FPCs online and taking low-priority FPCs offline.

- If power supply failure requires power management to power down some components, power management does so by gracefully powering down lower-priority FPCs.

Power management manages power to router components by employing a power budget policy. In its power budget policy, power management:

- Budgets power for each installed router component that requires power. The amount that power management budgets for each component is the maximum power that component might consume under worst-case operating conditions. For example, for the fan tray, power management budgets the amount of power required to run the fans at their maximum speed setting, even if the current fan speed is much lower.

- Manages the router for *N+N* power redundancy, which ensures uninterrupted system operation if one power supply fails.

- Provides power to host subsystem components, such as the Routing Engines, before it provides power to the FPCs.

- Manages the priority of individual FPCs. By assigning different priorities to the FPCs, you can determine which FPCs are more likely to receive power in the event of insufficient power.

Power Priority of FPCs

The power priority of FPCs determines:

- The order in which FPCs are allocated power.

- How power is reallocated if there is a change in power availability or demand in an operating router.

This section covers:

How an FPC's Power Priority Is Determined

Using the CLI, you can assign an explicit power priority to an FPC slot. The power priority is determined by the slot number, with the lowest-numbered slots receiving power first. Thus, if you do not explicitly assign priorities to slots, power priority is determined by slot number, with slot 0 having the highest priority. See Configuring Power-On Sequence to Redistribute the Available Power.

FPC Priority and FPC Power Allocation

When a PTX5000 is powered on, power management allocates power to components according to its power budget policy. After power management has allocated power to the host subsystem components, it allocates the remaining available power to the FPCs. It powers on the FPCs in the configured order of priority until all FPCs are powered on or the available power provided by both the PDUs is exhausted. Thus if available power is exhausted before all FPCs receive power, higher-priority FPCs are powered on while lower-priority FPCs remain powered off.

FPCs that have been taken offline are not allocated power.

> **NOTE**: Because power management does not allocate power to an FPC that has been taken offline, that FPC is brought online only when you commit a configuration. You must explicitly use the `request chassis fpc slot` *`slot-number`* `online` command to bring an FPC online that was taken offline previously.

If an FPC with a high priority in the priority sequence also has high-power requirement, and if the system does not have the required power available, then the lower priority FPCs with lower power requirements are also not powered on. This is to maintain consistency and also avoid powering off of the lower priority FPC when extra power is available. For example, if an FPC that requires 450 W has a higher priority than an FPC that requires 330 W, then the FPC with the lower power requirement (330 W) is also not powered on if the system does not have the required power to power the FPC that requires 450 W.

FPC Priority and Changes in the Power Budget

In an operating router, power management dynamically reallocates power in response to changes in power availability or demand or changes in FPC priority. Power management uses the configured priority on FPC slots to determine how to reallocate power in response to the following events:

- When a new power supply is brought online, FPCs that were powered off because of insufficient power are powered on in the order of priority.
- When a user changes the assigned power priority of one or more FPCs when power is insufficient to meet the power budget, power management reruns the current power budget policy and powers FPCs on or off based on their priority. As a result, FPCs receive power strictly by the order of priority and previously operating FPCs might no longer receive power.
- When an FPC is installed, Junos OS does not automatically power on and bring the FPC online. This FPC stays in the offline state until the user brings it online through the CLI or by pushing the online button, and only if the available chassis power is more than the budgeted power for this FPC, the FPC becomes operational.

Power Zones

In a PTX5000 equipped with high capacity PDUs and PSMs, there in one common zone that provides power to all FRUs and all FPCs. A high-capacity PDU can support up to eight PSMs and it does not support power zoning, unlike a normal-capacity PDU. All available PDU power is considered as a part of single zone. All PSMs provide power to the common zone. The PSM LEDs on the craft interface are interpreted as described in PTX5000 Craft Interface LEDs. After the PDU upgrade from the normal-capacity PDUs to High-Capacity PDUs, the power management converges all power zones into a single common zone. All FRU power is distributed based on the power available in the common zone.

NOTE: Presence of both normal-capacity PDUs and high-capacity PDUs is referred to as mixed-mode of operation and is supported only during the PDU upgrade.

To cater for the increase in the PIC power consumption, the power manager is enhanced to account for the PIC power separately from the FPC. The priority sequence for the PICs follows the priority sequence for the FPCs. That is, PICs installed in high-priority FPCs are given preference over PICs installed in low-priority FPCs. All PICs on an FPC have the same priority.

NOTE: You cannot mix existing PDUs with the High Capacity DC PDU.

Power Supply Redundancy

By default, power management in PTX5000 routers is configured to manage the power supplies for *N+N* redundancy, by which power supplies are held in reserve for backup if the other power supplies are removed or fail.

When power is insufficient to meet the budgeted power requirements, power management raises alarms as follows:

- With power supply redundancy, when one PSM fails, it does not cause FPCs to go offline. Only the `No redundant power supply alarm` is raised. However, with no redundancy, FPCs can go offline depending on the total chassis power available at that time. When an FPC or PIC goes offline due to insufficient power, which is indicated by `No power` in the output of the **show chassis fpc** command, then the `Insufficient Power - FRU(s) went offline` alarm is raised. The alarm gets cleared when there is sufficient power to bring up all the FPCs and PICs. The `Insufficient Power - FRU(s) went offline` alarm is raised when PSMs fail, when PSMs are powered off manually, or any time there is insufficient power for the system to power all the FPCs or PICs in the system.
- When power fails or when a PSM is removed, power management:
  - Calculates the total chassis power available from the remaining PSMs for the FPCs.
  - Powers off the FPCs based on the priority depending on the power budget for the FPCs and the FRUs and their configured power-on sequence.

> **NOTE**: In the scenario where the available power is more than the budgeted power required by the FPC but less than its maximum power, the FPC is taken offline and then brought online, but one or more PICs in that FPC are not online.

- When a new PSM is inserted, power management:
  - Checks the power-on sequence of the FPCs and the PICs and brings any offline PICs online when power is available.
  - Powers on the FPCs based on the FPC's budgeted power and its power-on sequence depending on its priority.
  - Maintains the power for high-priority FPCs and their PICs by taking the low-priority FPCs offline when all the FPCs are brought online, depending on the available power.

Power management clears all alarms when sufficient power is available to meet normal operating and reserved power requirements.

Power Redundancy on SRX5400

The power redundancy feature in SRX5400 supports to manage the high-capacity high line power supplies for 2+2 AC redundancy mode. The power rate is 1167W at low line and 2050W at high line on SRX5400. The 2+2 redundancy mode requires four AC power supplies.

The minimum PSU requirement is now 2 instead of 1 for the PEM alarm to be raised. If you install only 1 high-capacity high line AC, a major alarm is raised.

For more information about power supply on SRX5400 refer to SRX5400 Services Gateway AC Power Supply Specifications.

T4000 Power Management Overview

Starting with Junos OS Release 12.3, the power management feature is enabled on a Juniper Networks T4000 Core Router. This feature enables you to limit the overall chassis output power consumption. That is, this feature enables you to limit the router from powering on a Flexible PIC Concentrator (FPC) when sufficient output power is not available to power on the FPC during booting or normal operation.

The power management feature is enabled only when six input feeds with 40 amperes (A) each or four input feeds with 60 A each is configured on the router. The power management feature is *not* enabled for any other input feed–-current combination. When the power management feature is *not* enabled, Junos OS tries to power on all the FPCs connected to the router.

> **CAUTION**: : If you do not configure the power management feature and the maximum power draw is exceeded by the router during booting or normal operation, FPCs' states might change from Online to Offline or Present, some traffic might drop, or the interfaces might flap.

> **TIP**: Interface flapping occurs when a router alternately announces the state of the interface to be as *up* and *down* in quick sequence.

After you connect the input feeds to the router, you must configure the number of input feeds connected to the router and the amount of current received at the input feeds. Use the `feeds` statement and the `input current` statement at the `[edit chassis pem]` hierarchy level to configure the number of input feeds and the amount of current received at each input feeds, respectively.

> **NOTE**: You can connect three 80 A DC power cables to the six-input DC power supply by using terminal jumpers. When you do this, ensure that you configure the `feeds` statement to have the value `6` and the `input current` statement to have the value `40`. If these configurations are not set, the power management feature is *not* enabled and, therefore, Junos OS tries to power on all the FPCs connected to the router.

When the power management feature is enabled, FPCs connected to the router are powered on based on the power received by the router. If the router receives sufficient power to power on all the FPCs connected to the router, all the FPCs are powered on. If sufficient power is not available, Junos OS limits the number of FPCs brought online. That is, Junos OS uses the total available chassis output power as a factor to decide whether or not to power on an FPC connected to the router.

Of all the supported FPCs of a T4000 router, the T1600 Enhanced Scaling FPC4 (model number: T1600-FPC4-ES) has the greatest power requirement.Table 11 on page 120 compares the FPC connection limits between a six-input feed 40 A connection and a four-input feed 60 A connection when power management is enabled and T1600-FPC4-ES is connected to router.

Table 11: FPC Connection Limit Comparison

Six Input Feeds with 40 A Connection	Four Input Feeds with 60 A Connection
When T1600-FPC4-ES is *not* connected: • All eight FPC slots can be brought online.	When T1600-FPC4-ES is *not* connected: • A maximum of seven other FPCs can be brought online. That is, only seven slots out of the eight FPC slots can be brought online.
When only one T1600-FPC4-ES is connected: • A maximum of seven other FPCs can be brought online. That is, only seven slots out of the eight FPC slots can be brought online.	When only one T1600-FPC4-ES is connected: • A maximum of six other FPCs can be brought online. That is, only six slots out of the eight FPC slots can be brought online.
When only T1600-FPC4-ES FPCs are connected: • A maximum of six T1600-FPC4-ES FPCs can be brought online.	More than one T1600-FPC4-ES *cannot* be brought online.

NOTE:

- When the power management feature is enabled, FPC power-on consistency is not maintained across router reboots. That is, the same set of FPCs that were powered on before a reboot might not be powered on after the reboot. Before the router reboot, the FPCs are powered on according to their insertion order in the chassis. After the reboot, the FPCs are powered on according to the FRU power-on sequence configured in the `fru-poweron-sequence` statement at the `[edit chassis]` hierarchy level. If the FRU power-on sequence is not configured, Junos OS uses the ascending order of the slot numbers of the FPCs as the sequence to power on the FPCs.

- Removal of any online FPC from the chassis does not change the state of any other FPC and does not trigger the power management feature to power on the FPCs that were not powered on initially because of the lack of sufficient power. When any online FPC is removed from the chassis, if you need to trigger the power management feature to re-evaluate the situation, you need to reboot or restart the chassis. Alternatively, you can make a configuration change at the `[edit chassis]` hierarchy level and then issue the `commit` command to commit the changes made at the [edit chassis] hierarchy level. The power management feature to re-evaluates the situation when a configuration change is committed at the [edit chassis] hierarchy level.

Configuring the Six-Input DC Power Supply on T Series Routers

IN THIS SECTION

By default, the six-input DC power supply is configured to have all the six input feeds connected. You can also choose to provide four or five input feeds to the six-input DC power supply. When providing four or five input feeds on standalone routers, you need to configure the `feeds` statement at the `[edit chassis pem]` hierarchy level. When providing four or five input feeds to an LCC router in a routing matrix, you need to configure the `feeds` statement at the `[edit chassis lcc` *`lcc-number`* `pem]` hierarchy level.

Starting with Junos OS Release 12.3, the power management feature is enabled on T4000 routers with six-input DC power supply. The power management feature is enabled only when six input feeds with 40 amperes (A) each or four input feeds with 60 A each is configured on the router. To do this, you need to configure the `feeds` and `input-current` statements at the `[edit chassis pem]` hierarchy level.

NOTE:

- Before configuring input feeds for your router, see the *T640 Core Router Hardware Guide*, *T1600 Core Router Hardware Guide*, or *T4000 Core Router Hardware Guide* for special considerations and for the number of input feeds supported by the router.
- The value assigned to the `feeds` statement must be equal to the number of input feeds provided to the power supply. Else, an alarm message is generated to indicate the mismatch.

The following procedures describe how to configure the six-input DC power supply on different routers:

Configuring the Six-Input DC Power Supply on an LCC Router in a Routing Matrix

To configure the six-input DC power supply on an LCC router in a routing matrix:

1. At the `[edit chassis lcc lcc-number pem]` hierarchy level, configure the `feeds` statement with the number of input feeds provided to the power supply.

   ```
   [edit chassis lcc lcc-number pem]
   user@host# set feeds number-of-input-feeds
   ```

   For example:

   ```
   [edit chassis lcc 1 pem]
   user@host# set feeds 5
   ```

   NOTE: All power supplies in the router must use the same number of inputs feeds.

2. Verify the configuration by using the `show` command at the `[edit chassis]` hierarchy level:

   ```
   [edit chassis lcc 1 pem]
   user@host# show
   pem {
       feeds 5;
   }
   ```

Configuring the Six-Input DC Power Supply on T640 and T1600 Routers

To configure the six-input DC power supply on a standalone T640 or T1600 router:

1. At the `[edit chassis pem]` hierarchy level, configure the `feeds` statement with the number of input feeds provided to the power supply.

   ```
   [edit chassis pem]
   user@host# set feeds number-of-input-feeds
   ```

For example:

```
[edit chassis pem]
user@host# set feeds 5
```

NOTE: All power supplies in the router must use the same number of inputs feeds.

2. Verify the configuration by using the `show` command at the `[edit chassis]` hierarchy level:

```
[edit chassis]
user@host# show
pem {
    feeds 5;
}
```

Configuring the Six-Input DC Power Supply on T4000 Routers

To configure the six-input DC power supply on a T4000 router:

1. At the `[edit chassis pem]` hierarchy level, configure the `feeds` statement with the number of input feeds provided to the power supply.

```
[edit chassis pem]
user@host# set feeds number-of-input-feeds
```

For example:

```
[edit chassis pem]
user@host# set feeds 4
```

NOTE: All power supplies in the router must use the same number of inputs feeds.

2. Configure the input current received by the router.

```
[edit chassis pem]
user@host# set input-current amps-in-each-feed
```

For example, if the router receives 60 A of input current:

```
[edit chassis pem]
user@host# set input-current 60
```

> **NOTE**: You can connect three 80 A DC power cables to six-input DC power supply by using terminal jumpers. When you do this, ensure that you set the value of the `feeds` statement to `6` and that of the `input current` statement to `40`. If these configurations are not set, the power management feature is *not* enabled. For more information about the power management feature, see T4000 Power Management Overview.

3. Verify the configuration by using the `show` command at the `[edit chassis]` hierarchy level:

```
[edit chassis]
user@host# show
pem {
    feeds 4;
    input-current 60;
}
```

Redistributing the Available Power by Configuring Power-On Sequence

Routers running on Junos OS Release 10.0 and later support an enhanced AC Power Entry Module (PEM) to provide the necessary power infrastructure to support up to twelve higher-capacity DPCs with higher port density and slot capacity. To support the cooling requirements for the enhanced AC PEMs, the routers support enhanced fan trays and fans.

The default behavior for MPC power-on sequence is slot number based, that is, slot 0 is brought online first followed by slot 1, slot 2 up to slot 11. For the scenarios, where it is running a mix of high capacity line cards (for core facing), and low capacity line cards (for access facing) in their system, you can use the *fru-poweron-sequence* option to manually set the MPC power on sequence and hence ensure that the more important core facing line cards are brought online first irrespective of which slots these are in. This approach provides fine control over deterministically bringing up MPCs, howeverConfiguring Power-On Sequence to Redistribute the Available Powerr, it is heavy on configuration and entails to follow the discipline in slot to MPC mapping across all the systems.

The Junos OS enables you to configure the power-on sequence for the DPCs on an MX Series router chassis containing the new AC PEM. This enables you to redistribute the available power to the DPCs based on your requirements and the calculated power consumption of the DPCs. To configure the power-on sequence, refer to related information.

SEE ALSO

Configuring Power-On Sequence to Redistribute the Available Power

Configuring Power-On Sequence to Redistribute the Available Power

You can configure the power-on sequence for the Flexible PIC Concentrators (FPCs) on MX, PTX, and T routers. This configuration enables you to redistribute the available power to the FPCs on the basis of your requirements and the calculated power consumption of the FPCs.

To configure the power-on sequence:

1. At the `[edit chassis]` hierarchy level, configure the `fru-poweron-sequence` statement indicating the order in which the FPCs need to be powered on.

   ```
   [edit chassis]
   user@host# set fru-poweron-sequence fru-poweron-sequence
   ```

   For example:

   ```
   [edit chassis]
   user@host# set fru-poweron-sequence “0 2 1”
   ```

2. Verify the configuration by using the `show` command at the `[edit chassis]` hierarchy level:

   ```
   [edit chassis]
   ```

```
user@host# run show chassis power sequence
fru-poweron-sequence "0 2 1";
```

NOTE:

- If the configured sequence contains invalid numbers, Junos OS considers only the valid numbers in the sequence. The invalid numbers are silently discarded.
- If the power-on sequence is not configured by including the `fru-poweron-sequence` statement, Junos OS uses the ascending order of the slot numbers of the FPCs as the sequence to power on the FPCs.
- Issue the *show chassis power* command to view power limits and usage details for the FPCs.

SEE ALSO

fru-poweron-sequence

Configuring Voltage Level Monitoring of FPCs

IN THIS SECTION

You can monitor the voltage on the flexible PIC concentrator (FPC) at regular intervals. When the voltage falls below 10%, the FPC is offlined.

The faulty FPC is monitored at 500ms intervals. The output of the `show chassis fpc` command shows `Power Failure` for the faulty FPC. The FPC remains in powered down state until the voltage level is normal again.

Enabling Voltage Failure Errors on the FPC

`fpc-nmi-volt-fail-knob` controls the behavior of the FPC after detecting voltage failure, and to online or offline the FPC based on the voltage level. To enable monitoring the voltage level on the FPC:

1. Navigate to the `[edit chassis]` hierarchy level.

2. Include the `set chassis fpc-nmi-volt-fail-knob enable` statement to enable voltage monitoring on the FPC.

```
[edit chassis]
    {
        fpc-nmi-volt-fail-knob enable;
    }
```

Disabling Voltage Failure Errors on the FPC

To disable monitoring the voltage level on the FPC:

1. Navigate to the `[edit chassis]` hierarchy level.

2. Include the `set chassis fpc-nmi-volt-fail-knob disable` statement to disable voltage monitoring on the FPC.

```
[edit chassis]
    {
        fpc-nmi-volt-fail-knob disable;
    }
```

Overriding the Default Maximum Power (Junos OS Evolved)

IN THIS SECTION

On the PTX10001-36MR router, you can override the maximum power value of the power supply module (PSM) by specifying a lesser power value. Similarly, on the PTX10008 router, you can override the default power budget allocated to the line card by specifying a power value.

Overriding the Default Maximum Power (PTX10001-36MR)

You can override the maximum power value of a power supply module (PSM), if you need to deploy the PTX10001-36MR router in an environment that does not require the maximum power capacity (3000 W) of the PSM. You can use the command `set chassis psm max-power` to override the maximum power capacity of the PSM. Using this configuration, you can specify a value that is less than the maximum capacity of the PSM, and then monitor the real-time power consumption against the configured power value.

See the following example to know how to override the default power in PTX10001-36MR:

```
user@router# set chassis psm max-power 1600
user@router# commit
```

If the above configuration is set, the system power capacity is shown as 1600W. See the following `show chassis power detail` output:

```
user@router# show chassis power detail

Chassis Power        Voltage(V)    Power(W)

Total Input Power                    937
  PSM 0
    Input 1              229          391
    Output             12.03        305.44
    Capacity            1600 W (maximum 3000 W)
  PSM 1
    Input 1                0          546
    Output             12.04        515.08
    Capacity            1600 W (maximum 3000 W)

Item                 Used(W)
  Routing Engine 0       25
  CB 0                    5

System:
  Zone 0:
```

```
    Capacity:          3200 W (maximum 6000 W)
    Actual usage:      937 W
Total system capacity: 3200 W (maximum 6000 W)
```

NOTE: If the power consumption of the PTX10001-36MR router exceeds the threshold you configured using the `set chassis psm max-power` command, the software does not take any corrective action against the breach; and the router might still encounter a power failure.

If the power consumption exceeds the configured threshold, the system raises a chassis alarm, as shown in the following example:

```
user@router# show system alarms

Mar 15 12:51:30
2 alarms currently active
Alarm time              Class  Description
2020-03-15 12:50:52 UTC  Minor  Power consumption is critical
```

Overriding the Default Maximum Power (PTX10008)

On the PTX10008 router, during the system startup, the power management software by default takes the maximum power mentioned for each field replaceable unit (FRU) and makes the power calculations based on this number. However, you can override the default power budget allocated to the line card by specifying a power value (in watts). You can use the command `set chassis fpc` *`fpc-slot`* `max-power` *`watts`* to override the default power. You can use the command `show chassis fpc detail` to view the maximum power consumption by a line card.

You can also disable the power management on PTX10008 by using the command `set chassis no-power-budget`. If you disable the power management on PTX10008, the system does not move any of the FRUs to offline state in case of insufficient power. Instead, the system keeps all the FRUs powered on by default. However, in case of a power shortage, a power redundancy alarm is raised as shown in the following example.

```
user@router> show system alarms

1 alarm currently active

Alarm time Class Description
```

```
2019-07-25 21:16:25 UTC Major chassis No Redundant Powe
```

SEE ALSO

maximum-power

Powering Off Packet Forwarding Engines

You can power on or power off the Packet Forwarding Engines in a running system, or keep a Packet Forwarding Engine powered off when the FPC comes online. The following are a couple of scenarios in which this feature is used.

- When the Packet Forwarding Engine ASIC is malfunctioning.
- To conserve power in case the deployment does not require the full capacity of the system.

To power off a Packet Forwarding Engine, use the following steps:

`user@host#` **set chassis fpc *slot-number* pfe *pfe-id* power on**

`user@host#` **commit**

You need to apply this configuration to both the Packet Forwarding Engines in an ASIC to be able to commit the configuration.

> **NOTE**: On MX series routers with MPC10E-15C-MRATE, you can power off or power on only the Packet Forwarding Engine 2. The Packet Forwarding Engines 0 and 1 do not support this command. On the MPC10E-15C-MRATE, operating the Packet Forwarding Engine 2 requires the Packet Forwarding Engines 0 and 1 to be functional. You can use the command `show chassis fpc` *`fpc-lot`* `detail` to view the Packet Forwarding Engine power ON/OFF status and bandwidth for the individual Packet Forwarding Engines in the MPC10E-15C-MRATE.

You can use the show chassis fpc fpc-slot detail command to view the Packet Forwarding Engine power on/off configuration status. See an example below:

```
user@router> show chassis fpc 0 detail
Slot 0 information:
  State                              Online
```

```
  Temperature                      41 degrees C / 105 degrees F (PFE_24-HBM)
  Temperature                      44 degrees C / 111 degrees F (PFE_25-HBM)
  Temperature                      43 degrees C / 109 degrees F (PFE_26-HBM)
  Temperature                      41 degrees C / 105 degrees F (PFE_27-HBM)
  Temperature                      40 degrees C / 104 degrees F (PFE_28-HBM)
  Temperature                      40 degrees C / 104 degrees F (PFE_29-HBM)
  Temperature                      38 degrees C / 100 degrees F (PFE_30-HBM)
  Temperature                      39 degrees C / 102 degrees F (PFE_31-HBM)
  Start time                          2020-10-28 00:46:17 PDT
  Uptime                              1 day, 1 hour, 34 minutes, 48 seconds
  Max power consumption            825 Watts

PFE Information:

  PFE  Power ON/OFF  Bandwidth        SLC
  0    On            500
  1    On            500
  2    On            500
  3    On            500
  4    On            500
  5    On            500
  6    On            500
  7    On            500
```

Power Saving Mode (ACX7100-48L, ACX7100-32C)

IN THIS SECTION

Power saving mode enhances the energy efficiency of your ACX7100-48L and ACX7100-32C devices by selectively deactivating specific hardware components. Use unused ports to enable power saving mode and save power. Please note:

- After configuring this feature, you must reboot your system for the changes to take effect.
- You can configure this feature only on the unused ports.

Benefits of Power Saving Mode

- Energy Efficiency: By disabling certain hardware components and reducing the system's traffic handling capacity, Power Saving Mode can save approximately 40 watts of power, contributing to overall energy savings and lower operational costs.
- Extended Hardware Lifespan: Operating under reduced capacity can lead to less strain on hardware components, potentially extending the lifespan of the equipment and reducing the frequency of hardware replacements.
- Environmental Impact: Lower power consumption results in a smaller carbon footprint, supporting organizational sustainability goals and contributing to environmental conservation efforts.
- Customizable Performance: You have the flexibility to enable or disable Power Saving Mode based on your current network demands, allowing you to optimize your system performance and power usage according to your specific needs.

Enabling Power Saving Mode

You can enable the power saving mode by following the steps below:

1. At the `[edit chassis]` hierarchy level, use the `set interfaces interface-range powersaving` command to configure the power saving mode. When you know the unused ports range, use the following syntax:

   ```
   [edit chassis]
   user@host# set interfaces interface-range powersaving member-range member-range
   ```

   For example:

   ```
   user@host# set interfaces interface-range powersaving member-range et-0/0/24 to et-0/0/47
   ```

   To enable power saving mode on all the unused ports, use the following command:

   ```
   user@host# set interfaces interface-range powersaving unused
   ```

2. Reboot the system.

NOTE:

- If you try to configure power saving mode on an **unused** port with PTP enabled, the system will display an error message.
- The port range mentioned in the example is for ACX7100-48L.

Disabling Power Saving Mode

You can disable the power saving mode by following the steps below:

1. At the `[edit chassis]` hierarchy level, delete the configuration for ports or `member-range`. For example:

```
[edit chassis]
user@host# delete interfaces interface-range powersaving member-range et-0/0/24 to et-0/0/47
```

2. Reboot the system.

Use Feature Explorer to confirm platform and release support for specific features.

SEE ALSO

interfaces (QFX Series, ACX Series)

Low-Power Mode EM Policy Profile for Noise Reduction

SUMMARY

The Low-Power Mode Environment Management (EM) Policy Profile feature is designed to reduce the operational noise levels of MX10K4 and MX10K8 chassis when you use 100G ports for the LC9600 line card. By enabling this mode, you can lower the default minimum fan speed from 60% to 44%, addressing the need for quieter environments. We recommend to enable this feature only when you use 100G optics. This feature is particularly beneficial in acoustically sensitive premises such as data centers, where maintaining low noise levels is essential without compromising cooling efficiency.

To enable low-profile mode EM policy profile, use the following CLI command:

```
set chassis fpc-empolicy-profile low-power-mode
```

After enabling this configuration, you can use the `show chassis temperature-thresholds` or `show chassis fan` command to view the updated fan speed details.

SEE ALSO

No Link Title

No Link Title

Power Mode Management on PTX10002-36QDD

PTX10002-36QDD device supports two power supply units (PSUs) with 1 + 1 PSU redundancy. The operating mode of PTX10002-36QDD depends on the type (3000 W or 2200 W) of PSUs present in the system. Two 3000 W PSUs are required for the system to operate in normal power mode with 1 + 1 PSU redundancy. If one of the PSUs is absent, the system does not support 1 + 1 redundancy and determines the operating mode based on the available PSU. When you use a 2200 W PSU, the system operates in the low power mode.

You can also force the system to operate in low power mode when 3000 W PSUs are present, by using the following CLI command:

```
set chassis mode power-optimized
```

NOTE: You must reboot the system for the mode change to take effect.

This configuration reduces the power consumption of the device. See the following `show chassis power` command output:

```
user@host> show chassis power
Chassis Power          Voltage(V)    Power(W)

Total Input Power                      817
  PSM 0
    State: Online
    Input 1              51            444
    Output            12.03         427.17
```

```
  PSM 1
    State: Online
    Input 1              51          373
    Output            12.02        329.3

System:
  Power mode:             power-optimized
  Zone 0:
      Capacity:           6000 W (maximum 6000 W)
      Allocated power:    2150 W (3850 W remaining)
      Actual usage:       817 W
  Total system capacity: 6000 W (maximum 6000 W)
  Total remaining power: 3850 W
```

Power Redundancy for JNP10K-PWR-AC3 Power Supply Module

IN THIS SECTION

The JNP10K-PWR-AC3 power supply provides N+1 PSM (Power Supply Module) redundancy support for the MX10004, MX10008, PTX10004, PTX10008, and PTX10016 platforms. The JNP10K-PWR-AC3 PSM is equivalent to four power-supplies that consist of four input feeds (A0, A1, B0, and B1) with a maximum power output capacity of 7.8 KW. You can enable redundancy at the source or feed level.

The redundancy feature makes the system more reliable. It enables the system to raise an "alarm[1]" on page 139 when the power remaining in the system is less than the power of the connected PSM with the highest capacity.

When a PSM is removed or has a fault in a redundancy-disabled system, the power manager will power off the required FRUs.

See hardware guide for more information about JNP10K-PWR-AC3 PSM.

NOTE: JNP10K-PWR-AC3 PSMs cannot provide PSM redundancy in a 4 slot chassis with high power consuming line cards (LCs) such as JNP10K-LC1301.

NOTE: The configuration examples and alarms given below are for the MX10004 and MX10008 platforms. On the PTX10004, PTX10008, and PTX10016 platforms, use `psm` instead of `pem` when configuring the redundancy.

Source Redundancy

When you have two power sources (source A and source B), you can enable source redundancy by connecting two sets of independent power feeds from each source to the four input terminals of JNP10K-PWR-AC3 PSM.

Please note the following conditions before configuring the source redundancy:

- The software currently supports only one redundant source: for example, source A (main source) along with source B (backup or redundant source).
- The source redundancy is disabled by default. You can use the following CLI command to enable source redundancy:

```
set chassis pem redundancy source-redundancy
```

- Source redundancy is only applicable if all the PSMs in the system are JNP10K-PWR-AC3 PSMs. If you try to include a JNP10K-PWR-AC2 PSM in the power supply, the system will raise an "alarm[2]" on page 139.
- You must ensure that the feed distribution is even.
- It is not possible to enable source redundancy and feed redundancy simultaneously. Therefore, you need to disable feed redundancy before enabling source redundancy. To disable feed redundancy, use the following CLI command:

```
delete chassis pem redundancy feed-redundancy slot pem slot number
```

- You must set the DIP switches to the `feed expected` position for all four PSM input feeds; otherwise, the system will raise an "alarm[3]" on page 139.

Please note the following conditions after configuring the source redundancy feature:

- The Power Manager considers the capacity of each JNP10K-PWR-AC3 PSM as a 2-feed capacity. See hardware guide for more information.
- If one of the sources becomes unavailable, the system will raise an "alarm[4]" on page 139 corresponding to the failed feeds. In addition, the system will raise another "alarm[3]" on page 139 and disable the source redundancy until the source failure is fixed.
- The overall system power capacity reduces to protect against a source failure.
- The software will simulate the feature and determine the new system power capacity. If the system's new power capacity cannot accommodate the existing system load, the system will disable the source redundancy, raise an "alarm[5]" on page 139, and continue to operate in normal mode.
- When a PSM is faulty and there is no source failure, the system capacity is expected to reduce further. If the reduced system capacity cannot support the existing load, the system will disable the source redundancy, raise an "alarm[5]" on page 139, and continue to operate in normal mode. You should replace the failed PSM as soon as possible.

Feed Redundancy

You can enable feed rudundancy by connecting all the four input feeds (A0, A1, B0, and B1) of JNP10K-PWR-AC3 PSM to one or more power sources. When one feed is down, the other feed will continue providing the power and keep the platform operational.

> **NOTE**: Feed redundancy is not supported for JNP10K-PWR-AC2 PSMs and JNP10K-PWR-BLN3 active blank modules. On a Junos OS device, if you try to configure feed redundancy for these unsupported devices, the system will ignore the configuration and make a print available in the LCMD logs. On a Junos OS Evolved device, the system will make the `Feed redundancy unsupported for PSM` print available in the messages.

Please note the following conditions before configuring the feed redundancy:

- The PSM must have at least two feeds connected.
- The software currently supports only one redundant feed.
- You must set the DIP switches to the `feed expected` position for both primary and redundant feeds. Otherwise, the system will raise an "alarm[6]" on page 139. Additionally, you must confirm that the DIP switch configuration matches the connected feeds.

- Feed redundancy is disabled by default. You can use the following CLI command to enable feed redundancy:

```
set chassis pem redundancy feed-redundancy slot pem slot number
```

- It is not possible to enable source redundancy and feed redundancy simultaneously. Therefore, you need to disable source redundancy before enabling feed redundancy. To disable source redundancy, use the following CLI command:

```
delete chassis pem redundancy source-redundancy
```

Please note the following conditions after configuring the feed redundancy feature:

- The Power Manager will calculate the capacity of the JNP10K-PWR-AC3 PSM by subtracting one feed from the total number of connected feeds. See hardware guide for more information.
- If the redundant feed becomes unavailable, the system will raise an "alarm[4]" on page 139 corresponding to the failed feed. In addition, the system will raise another "alarm[7]" on page 139 and disable feed redundancy until the feed failure is fixed.
- The overall system power capacity reduces to protect against a feed failure.
- The software will simulate the redundancy feature and determine the system's new power capacity. If the system's new power capacity is not sufficient to accommodate the existing system load, the system will disable the feed redundancy on a Junos OS device. However, on a Junos OS Evolved device, the system will disable feed redundancy sequentially on PSMs until it achieves the required capacity. Both operating systems will also raise an "alarm[8]" on page 139 and continue to operate in normal mode.
- When a PSM is faulty and there is no feed failure, the system capacity is expected to reduce further. If the reduced system capacity is not sufficient to support the existing load, the system will disable the feed redundancy, raise an "alarm[8]" on page 139, and continue to operate in normal mode. You should replace the failed PSM as soon as possible.

Table 12: Power redundancy alarms

Sl No	Message displayed in the output of `show chassis alarms` command on a Junos OS device	Message displayed in the output of `show chassis alarms` command on a Junos OS Evolved device	Description
1	`No Redundancy`	`chassis No Redundant Power`	Appears when the power remaining in the system is less than the individual connected PSMs.
2	`PEM Source Redundancy Unsupported PEM`	`Unsupported PSM for source redundancy`	Appears when you connect an unsupported power supply.
3	`PEM Source Redundancy Failure`	`PSM Source Redundancy Failure`	Appears when you do not position the DIP switches for all feeds.
4	`PEM %d Feed` *`feed name`* `has no input source`	`PSM %d Input Feed` *`feed-name`* `Failed`	Appears when a source or redundant feed is not available.
5	`PEM Source Redundancy Unsupported`	`No Source Redundancy`	Appears when the power is not sufficient after enabling the source redundancy.
6	`PEM %d Feed Redundancy Expects Min 2 INP in DIP Switch Cfg`	`PSM %d Feed redundancy expects min 2 inputs in DIP switch config`	Appears when you do not position the DIP switches for all connected feeds.
7	`PEM %d FEED REDUNDANCY FAILURE`	`PSM %d Feed Redundancy Failure`	Appears when the feed redundancy is disabled.
8	`PEM Feed Redundancy Unsupported`	`PSM %d Feed Redundancy Unsupported`	Appears when the power is not sufficient after enabling the feed redundancy.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
18.2R1	Starting from Junos OS Release 18.2R1, for JNP10K-LC2101 MPC on MX10008 routers, dynamic power management is enabled by default.
17.3R1	Starting from Junos OS Release 17.3R1, for MX10003 routers, mic-aware dynamic power management is enabled by default.
17.2R1	Starting in Junos OS Release 17.2R1, EX9200 switches support dynamic power management.
17.2R1	In Junos OS Release 17.2R1, for EX9200 switches, dynamic power management for MICs is enabled by default.
15.1R1	Starting in Junos OS Release 15.1R1, MX Series routers support dynamic power management.
15.1R1	In Junos OS Release 15.1R1, for MX Series routers, dynamic power management for MICs is disabled by default.
15.1F5	In Junos OS Release 15.1F5 and later, dynamic power management is enabled by default on several MPCs.
14.1	Starting in Junos OS Release 14.1, the power management feature for PTX5000 routers ensures that at any time, the chassis power requirements do not exceed the available chassis power.

RELATED DOCUMENTATION

Configuring Ambient Temperature | 106

fru-poweron-sequence

5

CHAPTER

Managing Errors and Alarms

Resiliency

SUMMARY

This section covers generic information about the resiliency feature.

Resiliency represents the system's ability to anticipate, withstand, and rapidly recover from disruptions while maintaining critical functionality. This capability monitors the health status of various device components and handles faults by taking necessary actions.

Resiliency is a comprehensive solution applied at all levels of the system.

For hardware:

- Errors: ECC, AER, FEC, etc.
- Component failure: Control Board (CB) components (such as Memory, CPU, sensors, etc.), Packet Forwarding Engine (PFE) components (such as Fabric Links, PFE FPGA, etc.), and any Field Replaceable Units (FRUs) (such as Power Supply Modules, Fan Trays, etc.).
- Link stuck: Host path, CPU to Line card, PFE wedge, etc.

For software:

- Common error conditions, such as out of memory.
- Uncommon conditions.
- Corruptions, etc.

For protocols, when software and hardware fails, the protocol will choose to avoid that failed node.

When a hardware failure occurs, the resiliency module detects the anomaly and communicates to software for taking corrective actions, including but not limited to the following:

- Logs the message to give clear indication of failure details, including time stamp, module name, component name & failure details.
- Raises/clears alarms, if applicable.
- Glows the FRU fault if LED is present

- Performs local action, such as self-healing and taking the component out of service

You can use the following commands to view errors or faults on the device:

- show system errors
- show chassis errors active
- show chassis fpc errors
- show chassis led

NOTE: You can use the clear chassis fpc errors command to clear FPC errors on the device.

For further details about errors and alarms, see Understanding Chassis Alarms and Managing Errors. To configure alarms and errors at the component level, use CLI Reference Guide that lists all the associated CLI commands.

RELATED DOCUMENTATION

Understanding Chassis Alarms

Managing Errors

error

alarm

action (chassis error)

Understanding Chassis Alarms

SUMMARY

This topic lists various chassis conditions that are configured to trigger alarms. Chassis alarms are predefined alarms triggered by a physical condition on the device such as a power supply failure or excessive component temperature. You can use the `show chassis alarms` command to display the chassis

IN THIS SECTION

alarm information for presently active alarms. Chassis alarms are preset. You cannot modify them. You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm.

Chassis Conditions That Trigger Alarms

IN THIS SECTION

Various conditions related to the chassis components trigger yellow and red alarms. You cannot configure these conditions.

- "Backup Routing Engine Alarms" on page 185

- For PTX5000 Packet Transport Router chassis component alarm conditions, see the *PTX5000 Packet Transport Router Hardware Guide*
- For T320 Core Router chassis component alarm conditions, see the *T320 Core Router Hardware Guide*
- For T640 Core Router chassis component alarm conditions, see the *T640 Core Router Hardware Guide*
- For T1600 Core Router chassis component alarm conditions, see the *T1600 Core Router Hardware Guide*
- For T4000 Core Router chassis component alarm conditions, see the *T4000 Core Router Hardware Guide*
- For TX Matrix chassis component alarm conditions, see the *TX Matrix Router Hardware Guide*
- For TX Matrix Plus chassis component alarm conditions, see the *TX Matrix Plus Router Hardware Guide*

Chassis Component Alarm Conditions on M5 and M10 Routers

Table 13 on page 146 lists the alarms that the chassis components can generate on M5 and M10 routers.

Table 13: Chassis Component Alarm Conditions on M5 and M10 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at www.juniper.net/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace failed fan tray.	Red
Forwarding Engine Board (FEB)	The control board has failed. If this occurs, the board attempts to reboot.	Replace failed FEB.	Red
Flexible PIC Concentrator (FPC)	An FPC has failed. If this occurs, the FPC attempts to reboot. If the FEB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red

Table 13: Chassis Component Alarm Conditions on M5 and M10 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
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.	Replace failed component.	Red
Routing Engine	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red

Table 13: Chassis Component Alarm Conditions on M5 and M10 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Power supplies	A power supply has been removed from the chassis.	Install missing power supply.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red

Table 13: Chassis Component Alarm Conditions on M5 and M10 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The temperature sensor has failed.	Open a support case using the Case Manager link at www.juniper.net/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M7i and M10i Routers

Table 14 on page 149 lists the alarms that the chassis components can generate on M7i and M10i routers.

Table 14: Chassis Component Alarm Conditions on M7i and M10i Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow

Table 14: Chassis Component Alarm Conditions on M7i and M10i Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Compact FEB (CFEB)	For an M7i router, CFEB has failed. If this occurs, the board attempts to reboot.	Replace failed CFEB.	Red
	For an M10i router, both control boards have been removed or have failed.	Replace failed or missing CFEB.	Red
	Too many hard errors in CFEB memory.	Replace failed CFEB.	Red
	Too many soft errors in CFEB memory.	Replace failed CFEB.	Red
	A CFEB microcode download has failed.	Replace failed CFEB.	Red
Fan trays	A fan has failed.	Replace failed fan tray.	Red
	For an M7i router, a fan tray has been removed from the chassis.	Install missing fan tray.	Red
	For an M10i router, both fan trays are absent from the chassis.	Install missing fan tray.	Red
	For a TX Matrix Plus router, fan tray is not matching the ST-SIB-Ls SIB.	Install a Rev.3 fan tray.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router’s midplane from the front is broken.	Replace failed component.	Red

Table 14: Chassis Component Alarm Conditions on M7i and M10i Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Power supplies	A power supply has been removed.	Insert missing power supply.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
	For an M10i router, only one power supply is operating.	Insert or replace secondary power supply.	Red
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk. This alarm only applies, if you have an optional CompactFlash card.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red

Table 14: Chassis Component Alarm Conditions on M7i and M10i Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow

Table 14: Chassis Component Alarm Conditions on M7i and M10i Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M20 Routers

Table 15 on page 154 lists the alarms that the chassis components can generate on M20 routers.

Table 15: Chassis Component Alarm Conditions on M20 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below requires speed.	Replace fan tray.	Red
FPC	An FPC has failed. If this occurs, the FPC attempts to reboot. If the System and Switch Board (SSB) sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs in to the router's backplane from the front (generally, an FPC) is broken.	Replace failed component.	Red

Table 15: Chassis Component Alarm Conditions on M20 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red

Table 15: Chassis Component Alarm Conditions on M20 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
SSB	The control board has failed. If this occurs, the board attempts to reboot.	Replace failed control board.	Red
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow

Table 15: Chassis Component Alarm Conditions on M20 Routers ***(Continued)***

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M40 Routers

Table 16 on page 158 lists the alarms that the chassis components can generate on M40 routers.

Table 16: Chassis Component Alarm Conditions on M40 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filter	Change air filter.	Change air filter.	Yellow
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red
	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.	Replace failed FPC.	Red

Table 16: Chassis Component Alarm Conditions on M40 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	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.	Replace failed component.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply temperature sensor has failed.	Replace failed power supply or power entry module.	Yellow
	A power supply fan has failed.	Replace failed power supply fan.	Yellow
	A power supply has high temperature.	Replace failed power supply or power entry module.	Red
	A 5-V power supply has failed.	Replace failed power supply or power entry module.	Red
	A 3.3-V power supply has failed.	Replace failed power supply or power entry module.	Red
	A 2.5-V power supply has failed.	Replace failed power supply or power entry module.	Red

Table 16: Chassis Component Alarm Conditions on M40 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	A power supply input has failed.	Check power supply input connection.	Red
	A power supply has failed.	Replace failed power supply or power entry module.	Red
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow

Table 16: Chassis Component Alarm Conditions on M40 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
SCB	The System Control Board (SCB) has failed. If this occurs, the board attempts to reboot.	Replace failed SCB.	Red
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow

Table 16: Chassis Component Alarm Conditions on M40 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M40e and M160 Routers

Table 17 on page 163 lists the alarms that the chassis components can generate on M40e and M160 routers.

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filter	Change air filter.	Change air filter.	Yellow
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Connector Interface Panel (CIP)	A CIP is missing.	Insert CIP into empty slot.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the MCS sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	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.	Replace failed component.	Red
Miscellaneous Control Subsystem (MCS)	An MCS has an out of range or invalid temperature reading.	Replace failed MCS.	Yellow
	MCS0 has been removed.	Reinstall MCS0.	Yellow
	An MCS has failed.	Replace failed MCS.	Red
Packet Forwarding Engine Clock Generator (PCG)	A backup PCG is offline.	Set backup PCG online.	Yellow
	A PCG has an out of range or invalid temperature reading.	Replace failed PCG.	Yellow
	A PCG has been removed.	Insert PCG into empty slot.	Yellow
	A PCG has failed to come online.	Replace failed PCG.	Red

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Switching and Forwarding Module (SFM)	An SFM has an out of range or invalid temperature reading on SPP.	Replace failed SFM.	Yellow
	An SFM has an out of range or invalid temperature reading on SPR.	Replace failed SFM.	Yellow
	An SFM is offline.	Set SFM online.	Yellow
	An SFM has failed.	Replace failed SFM.	Red
	An SFM has been removed from the chassis.	Insert SFM into empty slot.	Red

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	All SFMs are offline or missing from the chassis.	Insert SFMs into empty slots or set all SFMs online.	Red
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red

Table 17: Chassis Component Alarm Conditions on M40e and M160 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M120 and M320 Routers

Table 18 on page 168 lists the alarms that the chassis components can generate on M120 routers.

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filters	Change air filter.	Change air filter.	Yellow
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Control Board (CB)	A CB Ethernet switch has failed.	Replace failed CB.	Yellow

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	A CB has been removed.	Insert CB into empty slot.	Red
	A CB has failed.	Replace failed CB.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
Forwarding Engine Boards (FEBs)	A spare FEB has failed.	Replace failed FEB.	Yellow
	A spare FEB has been removed.	Insert FEB into empty slot.	Yellow
	A FEB is offline.	Check FEB. Remove and reinsert the FEB. If this fails, replace failed FEB.	Yellow
	A FEB has failed.	Replace failed FEB.	Red
	A FEB has been removed.	Insert FEB into empty slot.	Red

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Host subsystem	A host subsystem has failed.	Replace the host subsystem.	Yellow
	A host subsystem has been removed.	Insert host subsystem into empty slot.	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.	Replace failed component.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has a high temperature.	Replace failed power supply or power entry module.	Red
	A power supply input has failed.	Check power supply input connection.	Red
	A power supply output has failed.	Check power supply output connection.	Red
	A power supply has failed.	Replace failed power supply.	Red

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	Chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red

Table 18: Chassis Component Alarm Conditions on M120 and M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on M320 Routers

Table 19 on page 174 lists the alarms that the chassis components can generate on M320 routers.

Table 19: Chassis Component Alarm Conditions on M320 Routers

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filters	Change air filter.	Change air filter.	Yellow
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Control Board (CB)	A CB has been removed.	Insert CB into empty slot.	Yellow

Table 19: Chassis Component Alarm Conditions on M320 Routers ***(Continued)***

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	A CB temperature sensor alarm has failed.	Replace failed CB.	Yellow
	A CB has failed.	Replace failed CB.	Red
CIP	A CIP is missing.	Insert CIP into empty slot.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the CB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red

Table 19: Chassis Component Alarm Conditions on M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	Too many soft errors in FPC memory.	Replace failed FPC.	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.	Replace failed component.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow

Table 19: Chassis Component Alarm Conditions on M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
	A spare SIB is missing.	Insert spare SIB in to empty slot.	Yellow
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
Switch Interface Board (SIB)	A SIB has failed.	Replace failed SIB.	Yellow

Table 19: Chassis Component Alarm Conditions on M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	A spare SIB has failed.	Replace failed SIB.	Yellow
	A SIB has an out of range or invalid temperature reading.	Replace failed SIB.	Yellow
	A SIB is missing.	Insert SIB into empty slot.	Red
	A SIB has failed.	Replace failed SIB.	Red
	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
Temperature	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red

Table 19: Chassis Component Alarm Conditions on M320 Routers *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	Chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms

Table 20 on page 179 lists the alarms that the chassis components can generate on MX Series 5G Universal Routing Platforms.

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filters	Change air filter.	Change air filter.	Yellow

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Dense Port Concentrators (DPC)s	A DPC is offline.	Check DPC. Remove and reinsert the DPC. If this fails, replace failed DPC.	Yellow
	A DPC has failed.	Replace failed DPC.	Red
	A DPC has been removed.	Insert DPC into empty slot.	Red
Fan trays	A fan tray has been removed from the chassis.	Install missing fan tray.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
	A higher-cooling capacity fan tray is required when an MPC is installed on the chassis.	Upgrade to a high-capacity fan tray.	Yellow
Host subsystem	A host subsystem has been removed.	Insert host subsystem into empty slot.	Yellow
	A host subsystem has failed.	Replace failed host subsystem.	Red

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
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.	Replace failed component.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has a high temperature.	Replace failed power supply or power entry module.	Red
	A power supply input has failed.	Check power supply input connection.	Red
	A power supply output has failed.	Check power supply output connection.	Red
	A power supply has failed.	Replace failed power supply.	Red
	Invalid AC power supply configuration.	When two AC power supplies are installed, insert one power supply into an odd-numbered slot and the other power supply into an even-numbered slot.	Red
	Invalid DC power supply configuration.	When two DC power supplies are installed, insert one power supply into an odd-numbered slot and the other power supply into an even-numbered slot.	Red

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	Mix of AC and DC power supplies.	Do not mix AC and DC power supplies. For DC power, remove the AC power supply. For AC power, remove the DC power supply.	Red
	Not enough power supplies.	Install an additional power supply.	Red
Routing Engine	Excessive framing errors on console port. An excessive framing error alarm is triggered when the default framing error threshold of 20 errors per second on a serial port is exceeded. This might be caused by a faulty serial console port cable connected to the device.	Replace the serial cable connected to the device. If the cable is replaced and no excessive framing errors are detected within 5 minutes from the last detected framing error, the alarm is cleared automatically.	Yellow
	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing CompactFlash card.	Reformat CompactFlash card and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default primary Routing Engine. If this fails, replace failed Routing Engine.	Yellow

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	System booted from hard disk.	Install bootable image on CompactFlash card. If this fails, replace failed Routing Engine.	Yellow
	CompactFlash card missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
	The Ethernet management interface (`fxp0` or `em0`) on the Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red
System Control Board (SCB)	An SCB has been removed.	Insert SCB into empty slot.	Yellow
	An SCB temperature sensor alarm has failed.	Replace failed SCB.	Yellow
	An SCB has failed.	Replace failed SCB.	Red

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C (131 degrees F), the fans have been turned on to full speed, and one or more fans have failed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and the fans have been turned on to full speed.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C (149 degrees F), and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red
	Chassis temperature has exceeded 75 degrees C (167 degrees F). If this condition persists for more than 4 minutes, the router shuts down.	• Check room temperature. • Check air filter and replace it. • Check airflow. • Check fan.	Red

Table 20: Chassis Component Alarm Conditions on MX Series 5G Universal Routing Platforms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
	The temperature sensor has failed.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red
Flexible PIC Concentrator (FPC)	FPC <slot number> Major Errors On MX Series routers with MPC1 and MPC2 line cards, a major chassis alarm is raised when the following transient hardware errors occur • CPQ Sram parity error • CPQ RLDRAM double bit ECC error By default, these errors result in the Packet Forwarding Engine interfaces on the FPC being disabled. You can use the `show chassis fpc errors` command to view the default or user-configured action that resulted from the error. You can check the syslog messages to know more about the errors.	To resolve the error, restart the line card. If the error is still not resolved, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Red

Backup Routing Engine Alarms

For routers with primary and backup Routing Engines, a primary Routing Engine can generate alarms for events that occur on a backup Routing Engine. Table 21 on page 186 lists chassis alarms generated for a backup Routing Engine.

NOTE: Because the failure occurs on the backup Routing Engine, alarm severity for some events (such as Ethernet interface failures) is yellow instead of red.

NOTE: For information about configuring redundant Routing Engines, see the Junos OS High Availability Library for Routing Devices.

Table 21: Backup Routing Engine Alarms

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The backup Routing Engine boots from an alternate boot device, the hard disk. The CompactFlash card is typically the primary boot device. The Routing Engine boots from the hard disk when the primary boot device fails.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Boot Device	The boot device (CompactFlash or hard disk) is missing in boot list on the backup Routing Engine.	Replace failed backup Routing Engine.	Red
Ethernet	The Ethernet management interface (`fxp0` or `em0`) on the backup Routing Engine is down.	• Check the interface cable connection. • Reboot the system. • If the alarm recurs, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Yellow

Table 21: Backup Routing Engine Alarms *(Continued)*

Chassis Component	Alarm Condition	Remedy	Alarm Severity
FRU Offline	The backup Routing Engine has stopped communicating with the master Routing Engine.	Open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).	Yellow
Hard Disk	Error in reading or writing hard disk on the backup Routing Engine.	Reformat hard disk and install bootable image. If this fails, replace failed backup Routing Engine.	Yellow
Multibit Memory ECC	The backup Routing Engine reports a multibit ECC error.	• Reboot the system with the board reset button on the backup Routing Engine. • If the alarm recurs, open a support case using the Case Manager link at www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Yellow

Chassis Component Alarm Conditions for Guest Network Functions (GNFs)

Table 22 on page 188 lists the Chassis conditions that trigger alarms on guest network functions (GNFs).

Read more about GNFs in this Junos Node Slicing article.

Table 22: GNF Alarms

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Mixed Master and Backup RE types This alarm is raised when the GNF primary Routing Engine and GNF Backup Routing Engine have been assigned either mismatching frequencies (with difference above 100 MHz), mismatching numbers of cores, or DRAM.	Correct the differences and then relaunch the corrected GNF Routing Engine.	Yellow
Routing Engine	System Incompatibility with BSYS The alarm is shown when any incompatibilities between BSYS and GNF software versions cause the GNF to go offline.	Make the required changes to the BSYS or GNF software through upgrade.	Red
Routing Engine	Feature Incompatibility with BSYS Indicates a minor incompatibility between BSYS and GNF software versions. This could result in a: • A warning error for the GNF. • A FRU going offline. **NOTE**: Minor incompatibilities do not cause the GNF to go offline.	Make the required changes to the BSYS or GNF software through upgrade.	Yellow

Chassis Component Alarm Conditions on SRX1500, SRX4100, SRX4200 and SRX4600 devices

Table 23 on page 189 lists the alarms that the chassis components can generate on SRX1500, SRX4100, SRX4200 and SRX4600 devices. Execute `show chassis alarms` operational mode command on SRX1500, SRX4100, SRX4200 and SRX4600 devices to view the alarm.

Table 23: Chassis Component Alarm Conditions on SRX1500, SRX4100, SRX4200 and SRX4600 devices

Chassis Component	Alarm Name/ Condition	Remedy	Alarm Severity
Power supply unit (PSU)	Appears when one among the two PSU is not available or not energized for SRX1500, SRX4100, and SRX4200.	Install the missing PSU or refer *pem abscence*	Red
Power supply unit (PSU)	Appears when one among the two PSU is not available or not energized for SRX4600.	Install the missing PSU or refer *pem abscence*	Yellow
FPC Line Card	**FPC Inefficient Port Mapping:** Appears when the two port blocks 0/0 - 0/3 and 0/4 - 0/7 are unequally used on the SRX4100 or SRX4200.	This minor alarm is triggered when the two port blocks 0/0 - 0/3 and 0/4 - 0/7 are unequally used. The alarm is cleared when the ports in UP status are more equally distributed over the two port blocks.	Yellow

MX204 LED Scheme Overview

LEDs on the interface cards display the status of the ports. In MX204 router, there are four port LEDs per port. Each port provides an individual status LED with four states signaled by the color/LED state: OFF, GREEN, AMBER, RED.

The following port LED display modes are defined:

- Normal—Represents the normal working mode of the LED. By default, the port status display mode is Normal.
- Port location—The port location mode is ON when a remote operator initiates a port location command for a port or a group of ports.

The following factors trigger a change in the port LED color:

- Change in the port state. For example, loss of signal (LOS) to no LOS, remote fault, or local fault
- Pluggable insertion or removal

- Change in configuration
- Activation or deactivation of port location feature

Table 24 on page 190 summarizes the state and color rules for the port LEDs. These rules help in determining the port LED color. When port location mode is activated, the port LED state or color can be determined from the Port Location ON column.

NOTE: In MX204 router, there are four port LEDs per port. On PIC 0, if the port operates at the speed of 40-Gbps or 100-Gbps, then the first LED of PIC 1 will be ON and the other three LEDs will be OFF. And, if the port operates at the speed of 10-Gbps, then all the LEDs will be ON.

Table 24: Port LED State and Color Rules

Pluggable Inserted	Explicitly Disabled	Port State	Normal	Port Location ON
Yes	No	Up	Green	Blinking green
Yes	No	Down; loss of signal (LOS) detected	Off	Blinking green
Yes	No	Down; transceiver hardware failure	Red	Blinking red
Yes	No	Down; any other fault other than LOS and transceiver hardware failure	Amber	Blinking amber
ANY	Yes	Port disabled by CLI	Amber	Blinking amber
No	No	Anything except disabled port; however, transceiver not present	Off	Blinking green

MPC and MIC Lane LED Scheme Overview

LEDs on the interface cards display the status of the ports. On some MICs and MPC that have multiple ports and supports multiple port speed, it is not feasible to have an individual LED display for each port on an interface card. Hence, a shared LED display is introduced—the lane LEDs.

The MX10003 MPC includes this new LED lane display. The Multi-Rate 12xQFSP28 MIC and the fixed-port PIC (6xQFSPP) have separate lane LEDs.

The lane LEDs of the MIC are located on the MIC itself, whereas the lane LEDs of the PIC are located on the MPC.

The following interface cards support lane LEDs:

- MX10003 MPC (Multi-Rate)
- Line card (MX10K-LC2101)
- Multi-Rate Ethernet MIC

You can select a port operating in a breakout mode for an individual lane display, either periodically or when the `request chassis port-led` command is executed. Similar to the port status LEDs, the lane LED supports 4 states defined by the color or the LED status—OFF, GREEN, AMBER, and RED.

Figure 1 on page 191 illustrates the port LED and lane LED displays on the MPC.

Figure 1: Port LED and Lane LED display on the MPC

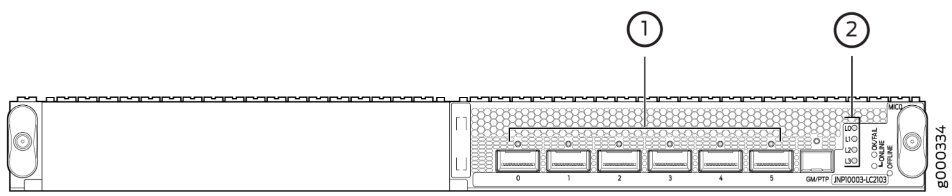

1– Port LEDs	2– Lane display LEDs

Figure 2 on page 191 illustrates the port LED and lane LED displays for the MPC.

Figure 2: Port LED and Lane LED display on the JNP10K-LC2101 MPC

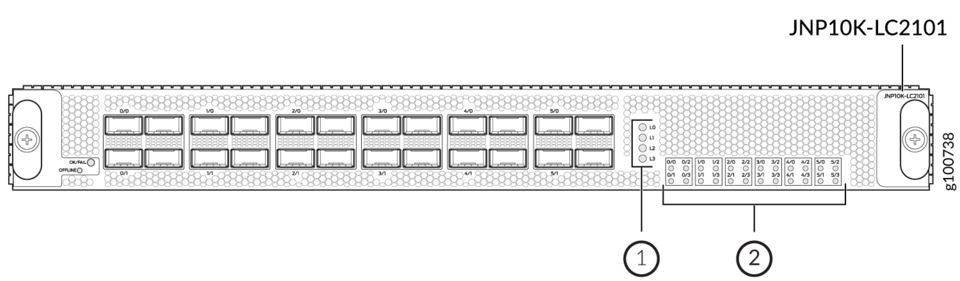

1– Lane display LEDs	2– Port LEDs

Figure 3 on page 192 illustrates the port LED and lane LED displays for the MIC.

Figure 3: Port LED and Lane LED display on the MIC

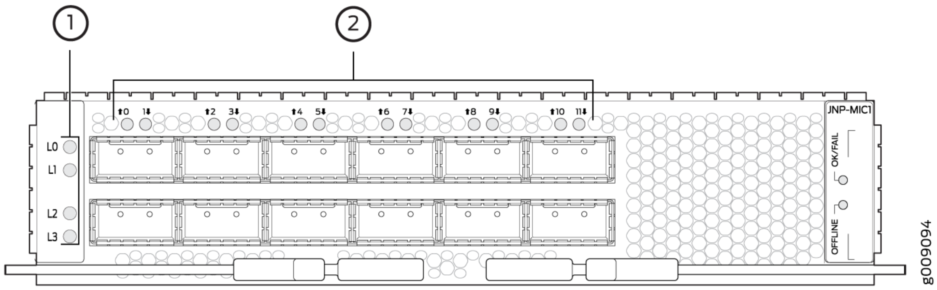

1– Lane display LEDs	2– Port LEDs

The following port LED display modes are defined:

- Normal—The port status LED represents port state or a breakout port state. By default, the port status display mode is Normal.
- Lane display—An array of lane status LEDs displays the status of each individual lane for the selected port. The lane display is ON when the software cycles through ports for lane status display. One port is selected at a time, and the display mode for that particular port switches to lane display mode. The other ports remain in normal display mode.
- Port location—The port location mode is ON when a remote operator initiates a port location command for a port or a group of ports. The `request chassis port-led` command temporarily overrides periodic software port selection for the lane display; all ports on an interface card that are not selected for port location switch to Normal mode, and selected ports switch to port location mode. If only one port is selected for port location, then the corresponding lane LEDs are applicable. However, if the selected port is in breakout mode, then all lane LEDs are applicable. If not in breakout mode, only lane 0 LED displays the port status. If more than one port is selected for port location, then the lane LEDs are disabled.

The following factors trigger a change in the port LED color:

- Change in the port state. For example, loss of signal (LOS) to no LOS, remote fault, or local fault
- Pluggable insertion or removal
- Change in configuration
- Activation or deactivation of port location feature
- Selection of breakout port for lane display

NOTE: Ports with all individual links in *Up* state are skipped and are not considered for lane display, thereby reducing the time needed to cycle through all the ports.

Table 25 on page 193 summarizes the state and color rules for the port LEDs. These rules help in determining the port LED color. When port location mode is activated, the port LED state or color can be determined from the Port Location ON column. If the breakout port is selected for the lane status display, then port LED state or color can be determined from the Lane Display column.

Table 25: Port LED State and Color Rules

Pluggable Inserted	Breakout Configuration State	Explicitly Disabled	Port State	Normal	Port Location ON	Lane Display
Yes	No breakout	No	Up	Green	Blinking green	-
Yes	No breakout	No	Down; loss of signal (LOS) detected	Off	Blinking green	-
Yes	No breakout	No	Down; transceiver hardware failure	Red	Blinking red	-
Yes	No breakout	No	Down; any other fault other than LOS and transceiver hardware failure	Amber	Blinking amber	-
ANY	No breakout	Yes	Port disabled by CLI	Amber	Blinking amber	-
No	Any	No	Anything except disabled port; however, transceiver not present	Off	Blinking green	-
Yes	Breakout	No	All breakout ports are UP	Green	Blinking green	Blinking green

Table 25: Port LED State and Color Rules *(Continued)*

Pluggable Inserted	Breakout Configuration State	Explicitly Disabled	Port State	Normal	Port Location ON	Lane Display
Yes	Breakout	No	All breakout ports are down with LOS	Off	Blinking green	Blinking green
Yes	Breakout	No	Hardware failure; transceiver initialization error at the port level (not individual lane)	Red	Blinking red	Blinking red
Yes	Breakout	Any	In all other cases the port LED color is amber	Amber	Blinking amber	Blinking amber

The following factors trigger a change in the lane LED color:

- A breakout port is selected for a lane display.
- Port location mode is activated for a port on a given interface card.

Table 26 on page 194 summarizes the state and color rules for the lane LEDs.

Table 26: Lane LED Color Rules

Pluggable Inserted	Breakout Configuration State	Explicitly Disabled	Port State	Order	LED Color
Yes	Breakout	No	Up	1	Green
Yes	Breakout	No	Down; loss of signal (LOS) detected	2	Off
Yes	Breakout	No	Down; transceiver hardware failure	3	Red

Table 26: Lane LED Color Rules *(Continued)*

Pluggable Inserted	Breakout Configuration State	Explicitly Disabled	Port State	Order	LED Color
Yes	Breakout	No	Down; fault other than LOS and transceiver hardware failure	4	Amber
Yes	Breakout	Yes	Breakout port is disabled in the CLI	5	Amber

Configuring Slow Packet Forwarding Engine Alarm

IN THIS SECTION

On an M Series, an MX Series, a T Series, or an SRX Series Firewall, the Packet Forwarding Engine might not send a resource acknowledgment message to the Routing Engine within a predetermined time of 360 seconds. This delay in receiving resource acknowledgment could be due to a slow or stuck Packet Forwarding Engine on the M Series, MX Series, T Series, or SRX Series Firewall, or on one of the LCCs connected to a TX Matrix, TX Matrix Plus, or TX Matrix Plus router with 3D SIBs.

Starting with Junos OS Release 13.2R1 (also applicable in Junos OS Releases 12.1R6, 12.2R5, 12.3R3, 13.1R2 and later), to display the issue as an alarm in the `show chassis alarms` command output and to append the alarm to the system log messages file, you must enable the slow Packet Forwarding Engine alarm on the router.

The following sections provide more information about the slow Packet Forwarding Engine alarm:

Enabling Slow Packet Forwarding Engine Alarm

To enable the slow Packet Forwarding Engine alarm, perform the following steps:

NOTE: By default, the slow Packet Forwarding Engine alarm is disabled.

1. In configuration mode, go to the `[edit chassis]` hierarchy level:

```
[edit]
user@host# edit chassis
```

2. Enable the slow Packet Forwarding Engine alarm by configuring the `slow-pfe-alarm` statement.

```
[edit chassis]
user@host# set slow-pfe-alarm
```

Disabling Slow Packet Forwarding Engine Alarm

To disable the slow Packet Forwarding Engine alarm, perform the following steps:

1. In configuration mode, go to the `[edit chassis]` hierarchy level:

```
[edit]
user@host# edit chassis
```

2. Disable the slow Packet Forwarding Engine alarm by deleting the `slow-pfe-alarm` statement.

```
[edit chassis]
user@host# delete slow-pfe-alarm
```

Verifying That the Alarm Output and System Log Messages are Updated

IN THIS SECTION

Purpose

To verify that the output of the `show chassis alarms` operational mode command and the system log messages file are updated with the slow Packet Forwarding Engine alarm when:

- The `slow-pfe-alarm` statement is enabled in the `[edit chassis]` hierarchy.
- The Packet Forwarding Engine resource acknowledgment is not received by the Routing Engine within a predetermined time of 360 seconds.

Action

To check the output on an M Series, MX Series, T Series, or an SRX Series Firewall:

1. Verify that the alarm is displayed in the output of the `show chassis alarms` operational mode command.

   show chassis alarms

   ```
   user@host> show chassis alarms
   1 alarms currently active
   Alarm time               Class  Description
   2013-02-05 01:12:33 PST  Minor  Potential slow peers are: XDPC2
   ```

   For field descriptions, see *show chassis alarms*.

2. Verify that the alarm is appended to the system log messages file.

   ```
   /var/log/messages -
   ... Alarm set: RE color=YELLOW, class=CHASSIS, reason=Potential slow peers are: XDPC2
   ... Minor alarm set, Potential slow peers are: XDPC2
   ```

To check the output on a TX Matrix, TX Matrix Plus, or a TX Matrix Plus with 3D SIBs router:

1. Verify that the alarm is displayed in the output of the `show chassis alarms` operational mode command.

   show chassis alarms

   ```
   user@scc> show chassis alarms
   scc-re0:
   ---------------------------------------------------------------------------
   9 alarms currently active
   Alarm time               Class  Description
   2013-02-06 00:45:46 PST  Minor  Potential slow peers are: LCC1 LCC0
   ```

```
...
lcc0-re0:
--------------------------------------------------------------------------
4 alarms currently active
Alarm time                Class  Description
2013-02-06 00:44:51 PST  Minor  Potential slow peers are: GFPC4 GFPC3
...
lcc1-re0:
--------------------------------------------------------------------------
4 alarms currently active
Alarm time                Class  Description
2013-02-06 00:45:44 PST  Minor  Potential slow peers are: GFPC10
...
lcc2-re0:
--------------------------------------------------------------------------
No alarms currently active
lcc3-re0:
--------------------------------------------------------------------------
No alarms currently active
```

For field descriptions, see *show chassis alarms*.

2. Verify that the alarm is appended to the system log messages file.

```
... Alarm set: RE color=YELLOW, class=CHASSIS, reason=Potential slow peers are: LCC0 LCC1
... Minor alarm set, Potential slow peers are: LCC0
... Alarm set: RE color=YELLOW, class=CHASSIS, reason=Potential slow peers are: GFPC4 GFPC3
... Minor alarm set, Potential slow peers are: GFPC4 GFPC3
... Alarm set: RE color=YELLOW, class=CHASSIS, reason=Potential slow peers are: GFPC10
... Minor alarm set, Potential slow peers are: GFPC10
```

Meaning

The output of `show chassis alarms` operational mode command and the system log messages file are updated as expected when the slow Packet Forwarding Engine alarm is enabled and when the Packet Forwarding Engine resource acknowledgment is not received by the Routing engine within a predetermined time of 360 seconds.

User-Defined Alarm Relay Overview

IN THIS SECTION

The ACX Series router alarm contact port—labeled ALARM on the front panel—allows you to manage sensors and external devices connected to the router in remote unstaffed facilities.

NOTE: Alarm contact port is not applicable on ACX5048 and ACX5096 routers.

Alarm Contact Port

The ACX Series router alarm contact port is a 15-pin D-type dry contact connector for alarms. The alarm contact port is used to generate LED alarms on the router and to turn external devices on or off. You can connect up to four input alarms and two output alarms. The alarm setting is open or closed.

Alarm Input

Alarm input provides dry contacts to connect to security sensors such as door or window monitors. The alarm input—open or closed—is sensed and reported to the management software. You can configure up to four alarm input relay ports (0 through 3) to operate as normally open or normally closed, and to trigger a red alarm condition or a yellow alarm condition or to ignore alarm conditions.

Alarm Output

Alarm output provides dry contacts to connect to external equipment, such as an audible or visual alarm that switches on or off–for example, a bell or a light. The four alarm output relay ports—0 through 3—are set up as follows:

- Ports 0 and 1—These ports can be configured to trigger an alarm when the system temperature goes to the red alarm status and when an alarm input port is triggered.
- Ports 2 and 3—These ports are *not* configured. They are used to indicate system major and minor alarms and are normally open. When a condition triggers an alarm, an alarm message is displayed.

To view the alarm input and output relay information, issue the `show chassis craft-interface` command from the Junos OS command line interface.

SEE ALSO

Configuring Chassis Alarm Relays
Configuring Chassis Alarm Input
Configuring Chassis Alarm Relays
relay (Chassis Alarm)

Configuring Chassis Alarm Relays

On ACX Series routers, you can configure alarm relays that can trigger alarms and turn external devices on or off. For example, if the router heats up to more than the critical temperature, the output port is activated and a device connected to the output port—such as a fan—is turned on.

To configure conditions that trigger alarms, include the `relay` statement with the `input` and `output` options at the `[edit chassis alarm]` hierarchy level.

```
[edit chassis alarm]
relay
    input {
        port port-number {
            mode (close | open);
            trigger (ignore | red | yellow);
        }
    }
    output{
        port port-number {
            input-relay input-relay;
            mode (close | open);
            temperature;
        }
    }
```

The following output shows an example configuration of a chassis relay alarm:

```
[edit chassis alarm]
user@host# show
relay {
    input {
        port 1 {
            mode close;
            trigger red;
        }
    }
    output {
        port 0 {
            temperature;
        }
    }
}
```

Configuring Chassis Alarm Input

The ACX Series router alarm contact port—labeled ALARM on the front panel—allows you to manage sensors and external devices connected to the router in remote unstaffed facilities. You can configure up to four alarm input ports (0 through 3) to operate as normally open or normally closed, and to trigger a red alarm condition or a yellow alarm condition or to ignore alarm conditions.

To configure an input alarm:

1. Configure the input port:

   ```
   [edit chassis alarm relay input port port-number]
   ```

   For example, to configure input port zero (0):

   ```
   user@host# edit chassis alarm relay input port 0
   ```

2. Configure the mode in which the input alarm is not active:

```
[edit chassis alarm relay input port port-number mode (close | open)]
```

For example, to configure open mode:

```
[edit chassis alarm relay input port 0]
user@host# set mode open
```

3. Configure the trigger to set off the alarm:

```
[edit chassis alarm relay input port port-number trigger (ignore | red | yellow)]
```

For example, to set off the yellow alarm:

```
[edit chassis alarm relay input port 0]
user@host# set trigger yellow
```

4. Verify the configuration with the `show` command:

```
[edit chassis alarm relay input port 0]
user@host# show
mode open;
trigger yellow;
```

5. Commit the configuration with the `commit` command.

To view the alarm input relay information, issue the `show chassis alarms` or `show chassis craft-interface` commands from the Junos OS command line interface.

Configuring Chassis Alarm Output

The ACX Series router alarm contact port—labeled ALARM on the front panel—allows you to manage sensors and external devices connected to the router in remote unstaffed facilities. You can configure up to two alarm output relay ports (0 and 1) to operate as normally open or normally closed, and to trigger an alarm when the system temperature goes to the red alarm status and when an alarm input port is triggered.

NOTE: Ports 2 and 3 are *not* configured. They are used to indicate system major and minor alarms and are normally open. When a condition triggers an alarm, an alarm message is displayed, and the corresponding LED turns on.

To configure an output alarm:

1. Configure the output port:

```
[edit chassis alarm relay output port port-number]
```

   For example, to configure output port zero (0):

```
user@host# edit chassis alarm relay output port 0
```

2. Configure the trigger to set off the alarm:

```
[edit chassis alarm relay output port port-number (input-relay | mode | temperature)]
```

   For example, to set off the alarm when the system temperature goes into the red status:

```
[edit chassis alarm relay output port 0]
user@host# set temperature
```

3. Verify the configuration with the `show` command:

```
[edit chassis alarm relay output port 0]
user@host# show
temperature;
```

4. Commit the configuration with the `commit` command.

To view the alarm output relay information, issue the `show chassis alarms` or `show chassis craft-interface` command from the Junos OS command line interface.

Configuring Chassis Alarm Input and Output (ACX710 Routers)

The alarm interface port, an RJ45 port on the front panel of the ACX710 router, provides user-configurable input and output signals. You can configure the alarm input to receive alarm inputs from the external devices (such as sensors) connected to the router through the alarm port. You can configure the alarm output to relay the alarms in the router to external alarm devices (for example, bells and bulbs) connected to the router through the alarm port. You can configure up to three alarm inputs and one alarm output.

The router supports configuration of up to three alarm inputs and one alarm output, using the command alarm-port at the [edit chassis] hierarchy. You can configure the alarm input signals independent of the alarm output signal, and vice versa.

Table 27: Alarm Port Pin-out Information

Pin Number of the Connector on the Device	Signal Definition	IN/OUT	CLI Mapping
1	ALARM_IN0_Sig	IN	port 1
2	ALARM_IN0_Return	IN	port 1
3	ALARM_IN1_Sig	IN	port 2
4	ALARM_IN2_Sig	IN	port 3
5	ALARM_IN1_Return	IN	port 2
6	ALARM_IN2_Return	IN	port 3
7	ALARM_OUT_Sig	OUT	port 1
8	ALARM_OUT_Return	OUT	port 1

To configure an alarm input:

1. Specify the input port number by using the command `set chassis alarm-port input port` *`port-number`*. The router supports three input ports (1 to 3).

   ```
   user@host# set chassis alarm-port input port 1
   ```

2. Configure a signal polarity for the alarm input based on the user environment.

   ```
   user@host# set chassis alarm-port input port 1 active low
   ```

3. Set the administrative state of the alarm input as enabled.

   ```
   user@host# set chassis alarm-port input port 1 admin-state enabled
   ```

4. Provide a description to the alarm input. For example, FAN.

   ```
   user@host# set chassis alarm-port input port 1 description FAN
   ```

5. Specify an alarm severity. The following are the available options: critical, major, minor, and warning.

   ```
   user@host# set chassis alarm-port input port 1 severity major
   ```

6. Commit the configuration with the `commit` command.

To view the input alarms, by using the `show chassis alarms` command.

To configure an alarm output:

1. Specify the output port number by using the command `set chassis alarm-port output port` *`port-number`*. The router supports only one output port (port number: 1).

   ```
   user@host# set chassis alarm-port output port 1
   ```

2. Set the administrative state of the alarm output as enabled.

   ```
   user@host# set chassis alarm-port output port 1 admin-state enabled
   ```

3. Provide a description to the alarm input.

```
user@host# set chassis alarm-port input port 1 description alarm-output-description
```

4. Commit the configuration with the `commit` command.

For more information, see *alarm-port*.

You can use the command `show chassis craft-interface` to view the alarm port configuration details.

```
user@router> show chassis craft-interface

System LED's on front panel:
-----------------------------
Fault LED :              On
Status LED :             Off
Operational LED :        On
Fan LED :                Off

Alarm-port on front panel:
-----------------------------

Input port :                1
         Active signal :    LOW
         Description :
         Admin state :      DISABLED
         Severity :         CRITICAL

Input port :                2
         Active signal :    LOW
         Description :
         Admin state :      DISABLED
         Severity :         CRITICAL

Input port :                3
         Active signal :    LOW
         Description :
         Admin state :      DISABLED
         Severity :         CRITICAL

Output port :               1
```

```
        Description :
        Admin state :       DISABLED
```

RELATED DOCUMENTATION

Silencing External Devices Connected to Alarm Relay Contacts | 225

Managing Errors

IN THIS SECTION

Configuring FPC Error Levels and Actions

Starting with Junos OS Release 13.3 or Release 14.2 for M320 routers, you can use MX Series, PTX Series, and T Series routers to configure Packet Forwarding Engine (PFE)-related error levels on FPCs and the actions to perform when a specified threshold is reached. In Junos OS Release 13.2 and earlier, Packet Forwarding Engine errors would disable the FPC. When you use the `error` command, Packet Forwarding Engine errors can be isolated, which reduces the need for a field replacement. Using the `error` command, you can classify errors according to severity, set an automatic recovery action for each severity, and configure the actions to perform when a specified threshold is reached. This command is available at the `[edit chassis fpc slot-number]` and `[edit chassis]` hierarchies.

To configure Packet Forwarding Engine error levels and actions for an FPC:

- (Optional) Configure the fatal error level threshold and action. A fatal error is an error that results in blockage of considerable amount of traffic across modules.

  ```
  [edit chassis fpc fpc-number error]
  user@host# set fatal action action
  user@host# set fatal threshold threshold-level
  ```

  If the severity level of the error is fatal, the action is carried out when the total number of errors reaches the threshold value. After the threshold value is crossed, for every occurrence of the error, an action is carried out.

- (Optional) Configure the major error level threshold and action. A major error is an error that results in continuing loss of packet traffic but does not affect other modules.

  ```
  [edit chassis fpc fpc-number error]
  user@host# set major action action
  user@host# set major threshold threshold-level
  ```

  If the severity level of the error is major, the action is carried out when the total number of errors reaches the threshold value. After the threshold value is crossed, for every occurrence of the error, an action is carried out.

- (Optional) Configure the minor error level threshold and action. A minor error is an error that results in the loss of a single packet but is fully recoverable.

  ```
  [edit chassis fpc fpc-number error]
  user@host# set minor action action
  user@host# set minor threshold threshold-level
  ```

  If the severity level is minor, the action is carried out only once when the total number of errors reaches the threshold value

Starting with Junos OS Release 18.1R3, MX Series routers support configuration of error thresholds and actions at the error scope and error category levels. Use the command `set chassis fpc` *`fpc-slot`* `error scope` *`error-scope`* `category` *`category`* `(fatal | major | minor) threshold` *`error-threshold`* `action (alarm | disable-pfe | get-state | offline | log | reset | trap | online-pfe | reset-pfe)` to configure a threshold and action for a particular error scope and category at the FPC level. You can also configure these features at the chassis level (at the [edit chassis] hierarchy). However, threshold and action configured at the [edit chassis fpc] hierarchy overrides the same configuration at the [edit chassis] hierarchy.

You can use the command show chassis fpc errors to view the error information at the error scope and category level.

For Junos OS Evolved, you can use the following show commands to view the error information:

- show system errors count—Displays system-wide errors and its count.
- show system errors active—Displays current active errors in the system.
- show system errors active fpc <slot number> —Displays active errors for the specified FPC.
- show system errors fru detail—Displays detailed FRU-specific error.
- show system errors fru detail fpc <slot number>—Displays information about detected errors based on the FRU.

If you have configured the action log against a particular error threshold, the system logs the event when the error count breaches the set threshold. The following sample syslog messages indicate an error threshold breach and the resultant action being taken:

```
Sep 17 23:12:10  sw-s3-u8-03 fpc0 Error: /fpc/0/pfe/0/cm/0/PE_Chip/1/
PECHIP_CMERROR_OQB_INT_REG_RD_ADDR_ERR (0x21078b), scope: pfe, category: functional, severity:
minor, module: PE Chip, type: Description for PECHIP_CMERROR_OQB_INT_REG_RD_ADDR_ERR
Sep 17 23:12:10  sw-s3-u8-03 fpc0 Performing action log for error /fpc/0/pfe/0/cm/0/PE_Chip/1/
PECHIP_CMERROR_OQB_INT_REG_RD_ADDR_ERR (0x21078b) in module: PE Chip with scope: pfe category:
functional level: minor
```

The offline, reset, disable-pfe, offline-pfe and reset-pfe actions are mutually exclusive with respect to configuration. The specified PFE is disabled automatically, if offline-pfe or reset-pfe is configured.

NOTE:
disable-pfe

The following table provides details about PFE error mapping actions and the system response:

Table 28: PFE Error Mapping Action and Response

Action	Response
disable-pfe	Disables all PFE interfaces, alarms and logs.

Table 28: PFE Error Mapping Action and Response ***(Continued)***

Action	Response
`offline`	Takes the FPC offline, disables the alarms and logs.
`reset`	Takes the FPC offline and resets to online, enables the alarms and logs.
`reset-pfe`	Powers-off the PFE, disables the alarms and logs, then, powers-on the PFE, enables the alarms and logs.
`offline-pfe`	Powers-off the PFE, disables the alarms and logs,

Example: Configuring FPC Error Detection and Self-Healing on T Series Core Routers

IN THIS SECTION

This example shows how to configure error detection and self-healing on a Juniper Networks T Series Core Router with Type 5 FPC.

Requirements

This example uses the following hardware and software components:

- Juniper Networks T4000 Core Router with Type 5 FPCs.
- Junos OS Release 13.3 or later.

Before you proceed, ensure that the required connections are complete and the interfaces are functional.

Overview

FPC error detection and self-healing involves configuring a set of actions to be performed on each FPC, when the number of errors for a particular severity increases beyond a user-configured threshold. The error severity is categorized into fatal, major, and minor. Recovery actions include raising an alarm, generating log entries, getting the current state of the FPC, restarting the FPC, taking the FPC offline, and resetting the FPC. For a particular FPC and error severity, you can configure the error threshold to any value within the allowed limits and map the threshold to an action. In this example, you will set these errors on FPC 0 in Juniper Networks T4000 Core Router.

Configuration

IN THIS SECTION

To configure the error detection and self-healing, you need to set the error severity, threshold values corresponding to each error severity, and actions to be performed when the threshold value is crossed.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit interfaces] hierarchy level.

```
set chassis fpc 0 fatal threshold 1 action resetset chassis fpc 0 major threshold 1 action
alarmset chassis fpc 0 minor threshold 10 action log
```

Configuring the Error Detection and Self-Healing

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the *Using the CLI Editor in Configuration Mode* and the CLI User Guide.

- Configure the threshold value and associated action for fatal errors.

  1. Set the error severity to fatal.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error fatal
     ```

  2. Set the threshold value for fatal errors.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error fatal threshold 1
     ```

  3. Set the associated action for fatal errors.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error fatal threshold 1 action reset
     ```

- Configure the threshold value and associated action for major errors.

  1. Set the error severity to major.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error major
     ```

  2. Set the threshold value for major errors.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error major threshold 1
     ```

  3. Set the associated action for major errors.

     ```
     [edit interfaces]
     ```

     ```
     user@host# set chassis fpc 0 error major threshold 1 action alarm
     ```

- Configure the threshold value and associated action for minor errors.

  1. Set the error severity to minor.

```
[edit interfaces]
```

```
[edit interfaces]
```

```
user@host# set chassis fpc 0 error minor
```

2. Set the threshold value for minor errors.

```
[edit interfaces]
```

```
user@host# set chassis fpc 0 error minor threshold 10
```

3. Set the associated action for minor errors.

```
[edit interfaces]
```

```
user@host# set chassis fpc 0 error minor threshold 10 action log
```

Results

The following is the result of the configuration for the fatal severity level.

```
user@host# set chassis fpc 0 error ?
Possible completions:
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
> fatal                FPC Fatal errors (default threshold = 1)
> major                FPC Major Level errors (default threshold = 1)
> minor                FPC Minor Level errors (default threshold = 10)user@host# set chassis fpc
0 error fatal action ?
Possible completions:
alarm              Raise FPC alarm
get-state          Retreive FPC state for debugging
log                Log occurence to system log file
offline            Offline FPC
offline-pic        Offline PICs associated with PFE on FPC
reset              Reset FPCuser@host# set chassis fpc 0 error fatal action resetuser@host# set
chassis fpc 0 error fatal threshold ?
Possible completions:
<threshold>        Error count at which to take the action (0..4294967295)user@host# set chassis
fpc 0 error fatal threshold 1
```

If you are done configuring the devices, enter `commit` from configuration mode.

Verification

IN THIS SECTION

To verify that the configuration is successful and the router in configured with the correct action, use the `show chassis fpc errors` command.

Verifying the Configured Actions Related to Fatal Severity of FPC Error

Purpose

Make sure that the threshold value and the associated action are set for fatal errors.

Action

```
user@host> show chassis fpc errors
FPC  Level Occurred Cleared Threshold Action-Taken Action
0    Fatal      0      0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED | pfe-4 -
ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
```

Meaning

The sample output shows `Fatal` error at FPC `0` with `0` error `Occurred` (no previous occurrences), `0` error `Cleared` (no previous occurrences) with `Threshold` value set to `1` and `Action-Taken` set to `RESET`.

Managing FPC Errors

IN THIS SECTION

On the PTX series routers, you can disable an FPC error or modify the severity of the error at the error-id level. See FPC self-healing for details on PTX platforms that support this feature.

The error-id, which uniquely identifies an FPC error, is represented in the uniform resource identifier (URI) format and is composed of a module identifier and an error identifier. If an error occurs, you can find the error-id in the system log messages.

Modifying Severity of an Error

Though you cannot configure a new error severity, you can modify the existing severity of an error. For example, if you do not want to treat a particular error (identified by an error-id) as fatal anymore, you can modify its severity to major or minor as required.

> **NOTE**: You cannot modify the error severity at a group (for example, category) level.

To modify the severity of an error, use the following command:

```
user@host# set chassis fpc fpc-slot error error-id severity new-severity
```

See the following example:

```
user@host# set chassis fpc 3 error "/cpu/0/memory/0/ECC_CORRECTED_ERROR"
severity minor
```

In the above example, you modified the severity of the error ID "`/cpu/0/memory/0/memory-uncorrected-error`" in FPC 3 to `minor`.

Disabling an Error

To configure the system to stop reporting an error, identify the error-id and disable it. You can find the error-id in the system log messages. To disable an error, use the following command:

```
user@host# set chassis fpc fpc-slot error error-id state disable
```

See the following example:

```
user@host# set chassis fpc 3 error "/cpu/0/memory/0/ECC_CORRECTED_ERROR"
state disable
```

In the above example, you disabled the error “`/cpu/0/memory/0/memory-uncorrected-error`” in FPC 3.

Powering Off Packet Forwarding Engines

You can power on or power off the Packet Forwarding Engines in a running system, or keep a Packet Forwarding Engine powered off when the FPC comes online. The following are a couple of scenarios in which this feature is used.

- When the Packet Forwarding Engine ASIC is malfunctioning.
- To conserve power in case the deployment does not require the full capacity of the system.

To power off a Packet Forwarding Engine, use the following steps:

```
user@host# set chassis fpc slot-number pfe pfe-id power off
user@host# commit
```

To power on a packet forwarding engine, use the following steps:

```
user@host# set chassis fpc slot-number pfe pfe-id power on
user@host# commit
```

NOTE: You need to apply this configuration to both the Packet Forwarding Engines in an ASIC to be able to commit the configuration.

NOTE: On MX series routers with MPC10E-15C-MRATE, you can power off or power on only the Packet Forwarding Engine 2. The Packet Forwarding Engines 0 and 1 do not support this command. On the MPC10E-15C-MRATE, operating the Packet Forwarding Engine 2 requires the Packet Forwarding Engines 0 and 1 to be functional. You can use the command `show chassis fpc` *`fpc-lot`* `detail` to view the Packet Forwarding Engine power ON/OFF status and bandwidth for the individual Packet Forwarding Engines in the MPC10E-15C-MRATE.

You can use the `show chassis fpc` *`fpc-slot`* `detail` command to view the Packet Forwarding Engine power on/off configuration status. See an example below:

```
user@router> show chassis fpc 0 detail

Slot 0 information:
  State                                 Online
  Temperature                        41 degrees C / 105 degrees F (PFE_24-HBM)
  Temperature                        44 degrees C / 111 degrees F (PFE_25-HBM)
  Temperature                        43 degrees C / 109 degrees F (PFE_26-HBM)
  Temperature                        41 degrees C / 105 degrees F (PFE_27-HBM)
  Temperature                        40 degrees C / 104 degrees F (PFE_28-HBM)
  Temperature                        40 degrees C / 104 degrees F (PFE_29-HBM)
  Temperature                        38 degrees C / 100 degrees F (PFE_30-HBM)
  Temperature                        39 degrees C / 102 degrees F (PFE_31-HBM)
  Start time                            2020-10-28 00:46:17 PDT
  Uptime                                1 day, 1 hour, 34 minutes, 48 seconds
  Max power consumption              825 Watts

PFE Information:

  PFE  Power ON/OFF  Bandwidth        SLC
  0    On            500
  1    On            500
  2    On            500
  3    On            500
  4    On            500
  5    On            500
  6    On            500
  7    On            500
```

Configuring Sanity Polling

You can configure the `sanity-poll` statement for a particular FPC or FEB or CFEB to start a periodic sanity check for that FPC or FEB or CFEB. The periodic sanity check includes checking for error conditions such as “register sanity issues,” “high temperature,” “hardware failure,” and so on. If you do not configure the `sanity-poll` statement, then sanity polling is disabled.

> **NOTE**: Currently, periodic sanity check is performed only on the routing chip register.

Sanity polling periodically checks for an error condition in an FPC or FEB or CFEB and performs the appropriate actions in case of an error.

- To configure sanity polling for an FPC on T Series routers and M320 routers, include the `sanity-poll` statement and its substatements at the `[edit chassis fpc slot-number]` hierarchy level:

```
[edit chassis]
fpc slot-number {
    sanity-poll {
        retry-count number;
        on-error {
            raise-alarm;
            power (cycle | off);
            write-coredump;
        }
    }
}
```

- To configure sanity polling for a FEB on the M120 router, include the `sanity-poll` statement and its substatements at the `[edit chassis feb slot-number]` hierarchy level:

```
[edit chassis]
feb slot-number {
    sanity-poll {
        retry-count number;
        on-error {
            raise-alarm;
            power (cycle | off);
            write-coredump;
        }
```

```
    }
}
```

- To configure sanity polling for a CFEB on M7i and M10 routers, include the `sanity-poll` statement and its substatements at the `[edit chassis cfeb slot-number]` hierarchy level:

```
[edit chassis]
cfeb slot-number {
    sanity-poll {
        retry-count number;
        on-error {
            raise-alarm;
            power (cycle | off);
            write-coredump;
        }
    }
}
```

NOTE: On a TX Matrix or TX Matrix Plus router, you can configure the `sanity-poll` statement at the `[edit chassis lcc number fpc number]` hierarchy level.

The `sanity-poll` statement comprises the following substatements:

- The `retry-count` statement specifies the number of rechecks to be performed after the occurrence of a particular error condition. If an error exists in all the periodic checks, then sanity polling reports an error and proceeds to perform the appropriate actions (described as options of the `on-error` statement).

  For example, if the periodic sanity check detects an error in the FPC or FEB or CFEB and if you configure the `retry count` *number* to 15, sanity polling does not report the error immediately. Sanity polling checks 15 times for the same error condition. If an error persists in all 15 rechecks, then it reports an error and takes the appropriate actions.

  If you do not configure the `retry-count` statement, then by default, the `sanity-poll` statement rechecks the detected error 10 times before reporting an error condition.

- If sanity polling detects an error condition, the `on-error` statement performs the appropriate actions to eliminate the error.

  The following actions are common to all kinds of error conditions:

- To generate a chassis alarm, configure the `raise-alarm` statement. The chassis alarm is displayed in the front panel of the chassis.
- To reboot the FPC or FEB or CFEB after generating a core file, configure the `power cycle` statement. This statement is useful for temporary software errors that are eliminated after reboot.
- To halt the FPC or FEB or CFEB, configure the `power off` statement. This statement is useful in case of permanent hardware failure.

  CAUTION: The `power off` statement halts the FPC. Ensure that you have backup paths through a different FPC or FEB or CFEB to avoid service outage.

  NOTE: The `power cycle` and `power off` statements are mutually exclusive: You can configure either the `power cycle` or the `power off` action for an error.

- To trigger the core file, configure the `write-coredump` statement.

You can configure multiple actions for a given FPC or FEB or CFEB. If you do not configure any actions, the `sanity-poll` statement generates only FPC or FEB or CFEB system log messages.

Configuring the Junos OS to Make a Flexible PIC Concentrator Stay Offline

By default, a Flexible PIC Concentrator (FPC) is configured to restart after a system reboot. You can use the `request chassis fpc` operational mode command to take an FPC offline, but on Junos OS the FPC attempts to restart when you enter a `commit` CLI command. To configure an FPC to stay offline and prevent it from restarting, include the `power off` statement at the `[edit chassis fpc` *`slot-number`*`]` hierarchy level:

```
[edit chassis fpc slot-number]
power off;
```

To bring an FPC online that is configured to stay offline and configure it to stay online, include the `power on` statement at the `[edit chassis fpc slot-number]` hierarchy level:

```
[edit chassis fpc slot-number]
power on;
```

Configuring an SFM to Stay Offline

By default, if you use the `request chassis sfm` CLI command to take a Switching and Forwarding Module (SFM) offline, the SFM attempts to restart when you enter a `commit` CLI command. To prevent a restart, you can configure an SFM to stay offline. This feature is useful for repair situations.

To configure an SFM to stay offline, include the `sfm` statement at the `[edit chassis]` hierarchy level:

```
[edit chassis]
sfm slot-number {
    power off;
}
```

- *slot number*—Slot number in which the SFM is installed.
- `power off`—Take the SFM offline and configure it to remain offline.

For example, the following statement takes an SFM in slot 3 offline:

```
[edit chassis]
sfm 3 power off;
```

Use the `show chassis sfm` CLI command to confirm the offline status:

```
user@host# show chassis sfm
                 Temp    CPU Utilization (%) Memory Utilization (%)
Slot    State    (C)     Total   Interrupt   DRAM (MB)   Heap   Buffer
0       Online   34      2       0           64          16     47
1       Online   38      2       0           64          16     47
2       Online   42      2       0           64          16     47
3       Offline --- Configured power off ---
```

To bring the SFM back online, delete the `edit chassis sfm` statement and then commit the configuration.

Resynchronizing FPC Sequence Numbers with Active FPCs when an FPC Comes Online

On M320, T320, T640, T1600, T4000, TX Matrix, and TX Matrix Plus routers, when you bring a Flexible PIC Concentrator (FPC) online, the sequence number on the FPC may not be synchronized with the other active FPCs in the router, which may result in the loss of a small amount of initial traffic.

To avoid any traffic loss, include the `fpc-resync` statement at the `[edit chassis]` hierarchy level. This ensures that the sequence numbers of the FPC that is brought online is resynchronized with the other active FPCs in the router.

```
[edit chassis]
fpc-resync;
```

NOTE: In order to prevent null-route filtering, the `fpc-resync` command will have no effect if a single LMNR based FPC and one or more I-chip FPCs exist in the same chassis.

Enabling a Routing Engine to Reboot on Hard Disk Errors

When a hard disk error occurs, a Routing Engine might enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding.

To recover from this situation, you can configure a single Routing Engine to reboot automatically when a hard disk error occurs. To enable this feature, include the `on-disk-failure reboot` statement at the `[edit chassis routing-engine]` hierarchy level.

```
[edit chassis routing-engine]
on-disk-failure {
    disk-failure-action (halt | reboot);
}
```

For dual Routing Engine environments, you can configure a backup Routing Engine to assume primary role automatically, if it detects a hard disk error on the primary Routing Engine. To enable this feature, include the `on-disk-failure` statement at the `[edit chassis redundancy failover]` hierarchy level. For information about this statement, see the Junos OS High Availability User Guide.

You can configure the Routing Engine to halt (instead of rebooting) when the hard disk fails on the Routing Engine. To configure this feature, include the `disk-failure-action (halt | reboot)` statement at the `[edit chassis routing-engine on-disk-failure]` hierarchy level:

```
[edit chassis routing-engine]
on-disk-failure {
    disk-failure-action (halt | reboot);
}
```

Use the **halt** option to configure the Routing Engine to halt when the hard disk fails. Use the **reboot** option to configure the Routing Engine to reboot when the hard disk fails.

Handling Thermal Health Events Using Thermal Health Check and PSM Watchdog

You can use the thermal health check feature to configure an action to be taken on detection of a thermal health event such as power leakage. The thermal check feature monitors the power supply module (PSM) power output and FRU power consumption and if it detects that the PSM power output exceeds the FRU power consumption by a user-defined threshold, it assumes that there is a thermal health event, and takes an action based on user configuration. You can configure actions such as auto shutdown or alarms to be initiated on detection of a thermal health event. An example of the configuration is as follows: `set chassis thermal-health-check action-onfail auto-shutdown shutdown-timer 10 power-threshold 700`. This example configuration enables the software to detect a thermal health event if the power leak exceeds 700W, and shuts down the system 10 seconds after the thermal health failure is detected.

The thermal health check feature works only if:

- The router has the high capacity AC or DC power distribution units (PDU) installed in both the slots, and each PDU has equal number of PSMs. Both AC PSM and DC PSM are supported.

  The supported PSMs and PDUs are listed below:

  - High Capacity AC PSM (model: PSM2-PTX-AC; firmware: 0210 or later; hardware revision: 06 or later)
  - High Capacity 60A DC PSM (model: PSM2-PTX-DC; firmware: 0315 or later; hardware revision: 09 or later)

  - High Capacity 60A DC PDU (model: PDU2-PTX-DC; use the firmware version 0404 or later with hardware revision 07; use the firmware version 0503 or later with hardware revision 08)
  - High Capacity AC Delta PDU (model: PDU2-PTX-AC-D; firmware: 0305 or later; hardware revision: 04 or later)
  - High Capacity AC Wye PDU (model: PDU2-PTX-AC-W; firmware: 0305 or later; hardware revision: 03 or later)
  - High Capacity Single Phase AC PDU (model: PDU2-PTX-AC-SP; firmware: 0102 or later; hardware revision: 03 or later)
- Each PDU has at least three PSMs that are online, and each online PSM is consuming above 60A current (in case of an AC PSM) or above 100A current (in case of a DC PSM).
- None of the FRUs (RE, SIB, and FPC) is in the 'Present' state.

On the router, you can also configure the PSM watchdog feature at the [edit chassis] hierarchy. If a thermal health event causes Junos to go down, the PSM watchdog feature detects it and shuts down the router. In the watchdog configuration, you can specify the watchdog timer in seconds. After the specified duration, the watchdog expires. You can also specify the frequency (in minutes) at which Junos resets the watchdog counter. If the watchdog counter doesn't get reset because of reasons such as Routing Engine crash, the PSM turns off the output power on watchdog timer expiry and thereby shuts down the router.

Example configurations are as follows:

- Use `set chassis psm watchdog timeout 600 pat-frequency 2`. This command enables PSM watchdog with the watchdog timer set to 600 seconds and the counter is set to be reset every 2 minutes.
- Use `set chassis thermal-health-check fet-failure-check action-onfail auto-shutdown shutdown-timer 10.`. This command enables thermal health check, and shutdowns the system, 10 seconds after FET failure is detected.

> **NOTE**: The PSM watchdog feature works only if all the online PSMs in the router support this feature.

In short, if the Routing Engine software is running when a thermal event occurs, the thermal health check feature detects the thermal event and takes an action. However, if the Routing Engine software goes down in a thermal health event, it is the PSM watchdog timer that detects this issue and brings down the system.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
13.3	Starting with Junos OS Release 13.3 or Release 14.2 for M320 routers, you can use MX Series, PTX Series, and T Series routers to configure Packet Forwarding Engine (PFE)-related error levels on FPCs and the actions to perform when a specified threshold is reached.

RELATED DOCUMENTATION

No Link Title

thermal-health-check

watch dog

Craft Interface

IN THIS SECTION

Silencing External Devices Connected to Alarm Relay Contacts

You can manually silence external devices connected to alarm relay contacts. To silence an external devices, press the alarm cutoff button located on the craft interface front panel of the device.

Silencing the device does not remove the alarm messages from the display (if present on the router or switch) or extinguish the alarm LEDs. In addition, new alarms that occur after an external device is silenced reactivate the external device.

Configuring the Junos OS to Disable the Physical Operation of the Craft Interface

You can disable the physical operation of the craft interface front panel on the router. When you disable the operation of the craft interface, the buttons on the front panel, such as the alarm cutoff button, no longer function. To disable the craft interface operation, include the `craft-lockout` statement at the `[edit chassis]` hierarchy level:

```
[edit chassis]
craft-lockout;
```

SEE ALSO

Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types

Silencing External Devices Connected to Alarm Relay Contacts | 225

Remote Port Identification using LEDs for Cabling Assistance

With new and higher-density Modular Interface Cards (MICs) and Modular Port Concentrators (MPCs), cabling is complex and can result in wiring mistakes. Remote port identification reduces the complexity by providing an easy way of identifying the ports that must be connected to the cables.Starting in Junos OS Release 16.1, the remote port identification feature is supported on MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E.

The remote port identification feature is supported on MX10K-LC2101 and MX10K-LC9600 on MX10008 routers.

LEDs, used to display the status of the port, can be configured to blink for a small duration of time to identify the port and provide cabling assistance. Depending on the port identification required, you can configure the LED of a specific port, LEDs of all ports, LED of a specific type of port to blink. For instance, on MX2020 routers with MPC8E, you can identify the active ports that support port speeds of 100 Gbps by configuring the LEDs of the specific port to blink. Similarly, you can identify active ports that support port speeds of 10 Gbps and 40 Gbps. You can configure the LED of, for example, active port 9 on the MX2020 router with MPC9E and MIC-MRATE. You can also make the LEDs of all the ports blink, if required.

You can specify the duration of time that a LED blinks. The default duration is 5 minutes (300 seconds). You can also stop the LED from blinking before the duration expires, if required.

To enable port identification on the enhanced MPCs, you can make the LED corresponding to the ports to blink using the `request chassis port-led` command.

Configuring the LCD Panel on EX Series Switches (CLI Procedure)

IN THIS SECTION

This topic applies to hardware devices in the EX Series product family, which includes switches and the XRE200 External Routing Engine, that support the LCD panel interface.

The LCD panel on the front panel of EX Series switches displays a variety of information about the switch in the Status menu and provides the Maintenance menu to enable you to perform basic operations such as initial setup and reboot. You can disable these menus or individual menu options if you do not want switch users to use them. You can also set a custom message that will be displayed on the panel.

Disabling or Enabling Menus and Menu Options on the LCD Panel

By default, the Maintenance menu, the Status menu, and the options in those menus in the LCD panel are enabled. Users can configure and troubleshoot the switch by using the Maintenance menu and view certain details about the switch by using the Status menu.

If you do not want users to be able to use those menus or some of the menu options, you can disable the menus or individual menu options. You can reenable the menus or menu options.

Issue the `show chassis lcd menu` operational mode command to see the menus or menu options that are currently enabled.

NOTE: On some platforms, you must specify an FPC slot number in these commands. See the lcd-menu statement for details.

To disable a menu:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-name disable
```

To enable a menu:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-name disable
```

To disable a menu option:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-option disable
```

To enable a menu option:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-option disable
```

Configuring a Custom Display Message

You can configure the second line of the LCD to display a custom message temporarily for 5 minutes or permanently.

To display a custom message temporarily:

- On an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4300 switch except EX4300-48MP and EX4300-48MP-S switches, a standalone EX4500 switch, a standalone EX4550 switch, an EX6200 switch, an EX8200 switch, or an XRE200 External Routing Engine:

  ```
  user@switch> set chassis display message message
  ```

- On an EX3300, EX4200, EX4300, EX4500, or EX4550 switch in a Virtual Chassis configuration:

  ```
  user@switch> set chassis display message message fpc-slot slot-number
  ```

To display a custom message permanently:

- On an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4300 switch except EX4300-48MP and EX4300-48MP-S switches, a standalone EX4500 switch, a standalone EX4550 switch, an EX6200 switch, an EX8200 switch, or an XRE200 External Routing Engine:

```
user@switch> set chassis display message message permanent
```

- On an EX3300, EX4200, EX4300 except EX4300-48MP and EX4300-48MP-S, EX4500, or EX4550 switch in a Virtual Chassis configuration:

```
user@switch> set chassis display message message fpc-slot slot-number permanent
```

NOTE: The buttons on the LCD panel are disabled when the LCD is configured to display a custom message.

To disable the display of the custom message:

```
user@switch> clear
chassis display message
```

You can view the custom message by issuing the `show chassis lcd` command.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
18.2	Starting in Junos OS Release 18.2, the remote port identification feature is supported on JNP10K-LC2101 on MX10008 routers.
16.1	Starting in Junos OS Release 16.1, the remote port identification feature is supported on MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E.

RELATED DOCUMENTATION

Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types *request chassis port-led*

6

CHAPTER

Network Services Mode

Configuring Network Services Mode

IN THIS SECTION

Network Services Mode Overview

IN THIS SECTION

A network services mode defines how the router chassis recognizes and uses certain modules. You can configure network services modes on MX Series 5G Universal Routing Platforms and on T4000 Core Routers with Type 5 FPCs.

On MX Series 5G Universal Routing Platforms, you can configure IP Network Services mode, Enhanced IP Network Services mode, Ethernet Network Services mode, or Enhanced Ethernet Network Services mode.

You can use either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode to improve the scaling and performance specific to filters in a subscriber access network that uses statically configured subscriber interfaces. For more information about using enhanced network services modes with firewall filters, see *Firewall Filters and Enhanced Network Services Mode Overview*.

On MX240, MX480, and MX960 routers, the MPC5E and MPC7E line cards power on only if the configured network services mode is `enhanced-ip` or `enhanced-ethernet`. All other MPCs work with any of the network services modes. MX2010 and MX2020 support only `enhanced-ip` and `enhanced-ethernet` network services modes.

NOTE: If Dense Port Concentrators (DPCs) in Ethernet Network Services mode or Enhanced Ethernet Network Services mode are up and running, you cannot configure the system for IP Network Services mode. You must first disable any Ethernet Network Services mode DPCs before switching to IP Network Services mode.

NOTE: When a chassis starts without any functioning FPCs, the Network Services mode defaults to IP Network Services. When the first FPC comes online, the configured Networks Services mode is applied.

NOTE: Starting from Junos OS Release 13.3, you can configure the Enhanced IP Network Services mode and Enhanced Ethernet Network Services mode on MX240, MX480 and MX960 routers with an SCBE2. Specify the `enhanced-ip` option or the `enhanced-ethernet` option at the `[edit chassis network-services]` hierarchy level.

You can configure T4000 Core Routers with Type 5 FPCs to run in Enhanced Network Services mode to enable improved virtual private LAN service (VPLS) MAC address learning. For more information, see *enhanced-mode*.

Table 29 on page 233 explains how different modules function when the MX Series 5G Universal Routing Platform chassis is configured to run in different network services modes.

Table 29: Network Services Mode Functions

Configuration Upon Boot or Configuration Change	Module Function
IP Network Services mode (default; upon boot)	All modules except DPCE-X and DPCE-X-Q are powered on. Starting with Junos OS Release 15.1, you can limit the maximum number of logical interfaces on MX Series routers with MS-DPCs to 64,000 for enhanced IP network services mode. To do this, include the `limited-ifl-scaling` option with the `network-services enhanced-ip` statement at the `[edit chassis]` hierarchy level. Using the `limited-ifl-scaling` option prevents any collision of logical interface indices that can occur in a scenario in which you enable the Enhanced IP Network Services mode on the router which also contains an MS-DPC.

Table 29: Network Services Mode Functions *(Continued)*

Configuration Upon Boot or Configuration Change	Module Function
Ethernet Network Services mode (upon boot)	All modules are powered on. However, operating in Ethernet Network Services mode restricts certain BGP protocol functions and does not support Layer 3 VPN, unicast RPF, and source and destination class usage (SCU and *DCU*) functions. In addition, the number of externally configured filter terms is restricted to 64K. Ethernet Network Services mode provides support for only Layer 2.5 functions.
Enhanced IP Network Services mode (upon boot)	Only MPCs, MS-MPCs, and MS-DPCs are powered on. **NOTE**: Only Multiservices DPCs (MS-DPCs) and MS-MPCs are powered on with the enhanced network services mode options. No other DPCs function with the enhanced network services mode options.
Enhanced Ethernet Network Services mode (upon boot)	Only MPCs, MS-MPCs, and MS-DPCs are powered on. All restrictions for operating in Ethernet Network Services mode apply. **NOTE**: Only Multiservices DPCs (MS-DPCs) and MS-MPCs are powered on with the enhanced network services mode options. No other DPCs function with the enhanced network services mode options.
Change from IP Network Services mode to Ethernet Network Services mode	DPCE-X and DPCE-X-Q modules are powered on. No reboot is required. No impact to MPCs or MS-DPCs.
Change from Ethernet Network Services mode to IP Network Services mode	Invalid modification. No commit occurs. A warning message indicates if any FPCs (along with their slot location) must be offline before switching to other network services. No impact to MPCs or MS-DPCs.
Change from Enhanced Ethernet Network Services mode to Enhanced IP Network Services mode.	Reboot is required.

Table 29: Network Services Mode Functions *(Continued)*

Configuration Upon Boot or Configuration Change	Module Function
Change from Enhanced IP Network Services mode to Enhanced Ethernet Network Services mode	Reboot is required.
Change from IP Network Services mode to Enhanced IP Network Services mode	System reboot is required (PFE/FPCs)
Change from Ethernet Network Services mode to Enhanced Ethernet Network Services mode	Reboot is required.

For details on the Layer 2.5 support for Ethernet Network Services mode, see " Restricted Software Features in Ethernet Network Services Mode" on page 237.

Network Services on SCBE2

The following scenarios are to be noted when you use an MX Series router with an SCBE2:

- You must configure the `set chassis network-services (enhanced-ip | enhanced-ethernet)` configuration command and reboot the router to bring up the FPCs on the router. However, after the router reboots, the MS DPC, the MX FPC, and the ADPC are powered off.
- All the FPCs and DPCs in the router are powered off when you reboot the router without configuring either the enhanced-ip option or the enhanced-ethernet option at the `[edit chassis network-services]` hierarchy level.
- You must reboot the router when you configure or delete the enhanced-ip option or the enhanced-ethernet option at the `[edit chassis network-services]` hierarchy level. The following warning message, which prompts you to reboot the router, is displayed when you configure or delete the enhanced-ip or the enhanced-ethernet configuration statement at the `[edit chassis network-services]` hierarchy level.

```
'chassis'
WARNING: Chassis configuration for network services has been changed. A
system reboot is mandatory. Please reboot the system NOW. Continuing without
a reboot might result in unexpected system behavior.
commit complete
```

- Starting with Junos OS Release 14.2, you must perform a commit synchronization of the settings between dual Routing Engines under some specific conditions. If you configure or remove the `enhanced-ip` or the `enhanced-ethernet` option at the `[edit chassis network-services]` hierarchy level on one of the Routing Engines on a router that contains dual Routing Engines, perform commit synchronization of the settings between the two Routing Engines by entering the `commit synchronize` command at the `[edit system]` hierarchy level. In addition, you must reboot all of the Routing Engines simultaneously (using the CLI command `request system reboot both-routing-engines`) when the enhanced IP network services mode is changed. The reboot is performed to prevent any unexpected system behavior.

NOTE: Dynamic multicast replication mode is supported on SCBE2. Static multicast replication mode is not supported on SCBE2.

NOTE: If a route's next hop is a unicast next hop through integrated routing and bridging (IRB) and the corresponding MAC address is learned over a label-switched interface (LSI), the IRB derives the Layer 2 information from the indirect next hop for the LSI. If you configure the `load-balance per-packet` policy statement, the indirect next hop of the LSI points to a unilist, which has all the member links to load balance the packets toward the MPLS cloud. You should configure the `enhanced-ip` option to enable the unicast next hop for IRB to use the unilist as the Layer 2 forwards next hop and load balance the packets.

Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers

You can configure MX Series 5G Universal Routing Platforms to run in different network services modes. Each network services mode defines how the chassis recognizes and uses certain modules.

To configure the network services mode of an MX Series router:

1. Access the chassis hierarchy.

```
[edit]
user@host# edit chassis
```

2. Specify the network services mode that you want the router to use.

```
[edit chassis]
user@host# set network-services service
```

Feature Restrictions on MX Series Routers Running in Ethernet Network Services Mode or Enhanced Ethernet Network Services Mode

Table 30 on page 237 lists Junos OS feature restrictions when running in Ethernet Network Services mode or Enhanced Ethernet Network Services mode.

Table 30: Restricted Software Features in Ethernet Network Services Mode

Software Feature	Restriction in Ethernet Network Services Mode
BGP	• Data plane support applies only to Ethernet and MPLS. • BGP only supports the following address families: inet labeled-unicast, inet unicast, inet-vpn unicast, l2vpn, and route-target.
L3VPN	Layer 3 VPNs are supported. You can only include loopback interfaces in the Virtual Routing and Forwarding (VRF) instance. A maximum of two VRFs are supported. Each VRF can handle up to 10,000 routes. The `ping mpls l3vpn` operational mode command is also supported.
Unicast RPF	Unicast reverse-path forwarding is disabled.
Source and destination class usage (SCU and DCU)	Source and Destination Class Usage is disabled.
Filter terms	The number of externally configured filter terms is restricted to 64 KB.
Prefixes	The number of supported prefixes is restricted to 32 K.

NOTE: MX Series routers supporting Layer 2.5 functions work as full-scale routers and they support interior gateway protocol (IGP), multicast routing protocols, and other routing features. The restrictions applicable on these routers are that the number of routes is limited and you cannot use BGP.

Limiting the Maximum Number of Logical Interfaces on MX Series Routers With MS-DPCs in Enhanced IP Network Services Mode

Starting in Junos OS Release 15.1, you can impose a limitation on the maximum number of logical interfaces on MX Series routers with MS-DPCs to be 64,000 for enhanced IP network services mode. To impose that limit, include the `limited-ifl-scaling` option with the `network-services enhanced-ip` statement at the `[edit chassis]` hierarchy level. When network-services is configured as enhanced IP mode, the kernel increases the total number of logical interfaces to 256,000. However, MS-DPC line cards are not capable of handling more than 64,000 logical interfaces globally on a router. Using the `limited-ifl-scaling` option prevents the problem of a collision of logical interface indices that can occur in a scenario in which you enable enhanced IP services mode and an MS-DPC is also present in the same chassis. To support MS-DPCs with enhanced IP mode on the chassis, you must limit the maximum logical interfaces as 64,000, which is performed with the `limited-ifl-scaling` option.

To define the maximum number of logical interfaces on MX Series routers with MS-DPCs as 64,000, include the `limited-ifl-scaling` option with the `network-services enhanced-ip` statement at the `[edit chassis]` hierarchy level.

```
[edit chassis]
network-services enhanced-ip limited-ifl-scaling;
```

When the default network services mode on a router is IP services mode (by using the `network-services ip` statement), the maximum logical interfaces is set as 64,000. When you change the network services mode as enhanced IP, the chassis process sets a general configuration (GENCFG) script to the kernel that increases the maximum logical interfaces as 256,000. When you configure the `limited-ifl-scaling` option with the `network-services enhanced-ip` statement, the chassis process does not generate a message to the kernel to increase the number of logical interfaces. As a result, the kernel retains the maximum number of logical interfaces as 64,000.

If your router chassis is previously configured with enhanced IP services mode and without the `limited-ifl-scaling` option set, and if you later configure the setting to limit the logical interfaces for MS-DPCs, the number of logical interfaces remains as 256,000 and it is not reduced. A cold reboot of the router must be performed in such a case to reduce the logical interfaces after you set the `limited-ifl-scaling` option with the `network-services enhanced-ip` statement. When you enter the `limited-ifl-scaling` option, none of the MPCs are moved to the offline state. All the optimization and scaling capabilities supported with enhanced IP mode apply to enhanced IP mode with the limitation of IFL scaling functionality.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.1	Starting with Junos OS Release 15.1, you can limit the maximum number of logical interfaces on MX Series routers with MS-DPCs to 64,000 for enhanced IP network services mode.
15.1	Starting in Junos OS Release 15.1, you can impose a limitation on the maximum number of logical interfaces on MX Series routers with MS-DPCs to be 64,000 for enhanced IP network services mode.
14.2	Starting with Junos OS Release 14.2, you must perform a commit synchronization of the settings between dual Routing Engines under some specific conditions.
13.3	Starting from Junos OS Release 13.3, you can configure the Enhanced IP Network Services mode and Enhanced Ethernet Network Services mode on MX240, MX480 and MX960 routers with an SCBE2.

RELATED DOCUMENTATION

Firewall Filters and Enhanced Network Services Mode Overview

Junos OS Subscriber Management and Services Library

Configuring Enhanced IP Network Services for a Virtual Chassis

enhanced-mode

network-services

7

CHAPTER

Packet Scheduling Mode

Enabling an M160 Router to Operate in Packet Scheduling Mode

By default, packet scheduling is disabled on M160 Routers. 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 application-specific integrated circuit (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 SFMs and FPCs and restarts them in the new mode.

NOTE: Packet scheduling is available for M160 routers only.

RELATED DOCUMENTATION

Configuring the Junos OS to Support Eight Queues on IQ Interfaces for T Series and M320 Routers

8

CHAPTER

OSS Mapping

Configuring OSS Mapping

IN THIS SECTION

Configuring OSS Mapping to Represent a T4000 Chassis as a T1600 or a T640 Chassis

IN THIS SECTION

You can configure the operations support systems (OSS) mapping feature to represent a T4000 chassis as a T1600 chassis or a T640 chassis. This topic includes the following tasks:

Configuring T4000 Chassis as a T1600 Chassis

To configure a T4000 chassis as a T1600 chassis:

1. In configuration mode, go to the `[edit chassis]` hierarchy level.

```
[edit]
user@T4000# edit chassis
```

2. Configure the OSS mapping feature to map the T4000 chassis to a T1600 chassis.

```
[edit chassis]
user@T4000# set oss-map model-name t1600
```

Configuring T4000 Chassis as a T640 Chassis

To configure a T4000 chassis as a T640 chassis:

1. In configuration mode, go to the `[edit chassis]` hierarchy level.

```
[edit]
user@T4000# edit chassis
```

2. Configure the OSS mapping feature to map the T4000 chassis to a T640 chassis.

```
[edit chassis]
user@T4000# set oss-map model-name t640
```

NOTE: By default, the OSS mapping feature is disabled.

Disabling the OSS Mapping Feature

To disable the OSS mapping feature:

1. In configuration mode, go to the `[edit chassis]` hierarchy level.

```
[edit]
user@T4000# edit chassis
```

2. Disable the OSS mapping feature that maps a T4000 chassis to a T640 chassis.

```
[edit chassis]
user@T4000# delete oss-map model-name t640
```

3. Disable the OSS mapping feature that maps a T4000 chassis to a T1600 chassis.

```
[edit chassis]
user@T4000# delete oss-map model-name t1600
```

NOTE:

- The `set chassis oss-map model-name t640 | t1600` command is applicable only on T4000 routers. You must explicitly set this command when a T1600 chassis or a T640 chassis is upgraded to a T4000 chassis.
- You can execute the `set chassis oss-map model-name t640` command or the `set chassis oss-map model-name t1600` command if the OSS is compatible with either the T640 chassis or the T1600 chassis, respectively.

Example: Configuring a T4000 Chassis to Represent a T640 Chassis

IN THIS SECTION

This example shows how to configure OSS mapping feature to represent a T4000 chassis as a T640 chassis. You can extend this concept to configure a T4000 chassis to represent as a T1600 chassis as well.

Requirements

This example uses the following hardware and software components:

- One T4000 router
- Junos OS Release 12.3R3, 13.1R2, 13.2R1, or later

Overview

Operations support systems (OSS) is used by service providers to maintain their networks. When a new router is added or removed from the network, the OSS must be updated to reflect the changes. This process is tedious and time-consuming.

When a T1600 chassis or a T640 chassis is upgraded to a T4000 chassis, the OSS identifies the new chassis as a new networking element and follows a tedious process of qualifying it for the customer's network. The *OSS mapping feature* helps avoid this scenario.

Using the OSS mapping feature, you can map a T4000 chassis to a T1600 chassis or to a T640 chassis with the `set chassis oss-map model-name t640|t1600` configuration command. This configuration command overrides the chassis model name, so that the OSS recognizes the chassis as a known chassis and proceeds without any requalification.

NOTE:

- The `set chassis oss-map model-name t640 | t1600` command is applicable only on T4000 routers. You must explicitly set this command when a T1600 chassis or a T640 chassis is upgraded to a T4000 chassis.
- You can execute the `set chassis oss-map model-name t640` command or the `set chassis oss-map model-name t1600` command, if the OSS is compatible with either the T640 chassis or the T1600 chassis, respectively.

Configuring the T4000 Chassis to Represent a T640 Chassis

IN THIS SECTION

Procedure

Step-by-Step Procedure

To configure the T4000 chassis to represent a T640 chassis by using the OSS mapping feature:

1. In configuration mode, go to the `[edit chassis]` hierarchy level.

```
[edit]
user@T4000# edit chassis
```

2. Configure the OSS mapping feature to map the T4000 chassis to a T640 chassis.

```
[edit chassis]
user@T4000# set oss-map model-name t640
```

Verification

IN THIS SECTION

Verifying the OSS Mapping Feature

Purpose

To verify that the OSS mapping feature is working on a T4000 router.

Action

Run the show chassis operational command and verify that the configured known chassis name is displayed when the T4000 chassis is mapped to a T640 chassis.

- Run the `show chassis hardware` operational command:

  show chassis hardware

```
user@T4000> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11B3892AHA      T640
Midplane         REV 01   710-027486   RC9848            T-series Backplane
```

```
FPM GBUS         REV 13   710-002901   BBAG5143         T640 FPM Board
FPM Display      REV 04   710-021387   BBAL2705         T1600 FPM Display
CIP              REV 06   710-002895   BBAL3705         T-series CIP
PEM 1            REV 03   740-036442   VJ00054          Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAJ0727         T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAE3887         T640 Sonet Clock Gen.
Routing Engine 0 REV 06   740-026941   P737F-002705     RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002675     RE-DUO-1800
CB 0             REV 09   710-022597   EF7371           LCC Control Board
....
```

- Run the `show chassis hardware detail` operational command:

show chassis hardware detail

```
user@T4000> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN11B3892AHA     T640
Midplane         REV 01   710-027486   RC9848           T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5143         T640 FPM Board
FPM Display      REV 04   710-021387   BBAL2705         T1600 FPM Display
CIP              REV 06   710-002895   BBAL3705         T-series CIP
PEM 1            REV 03   740-036442   VJ00054          Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAJ0727         T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAE3887         T640 Sonet Clock Gen.
Routing Engine 0 REV 06   740-026941   P737F-002705     RE-DUO-1800
  ad0     3823 MB  SMART CF            201101050335CCFACCFA Compact Flash
  ad1    62720 MB  SMART Lite SATA Drive 2011021700D8789F789F Disk 1
Routing Engine 1 REV 06   740-026941   P737F-002675     RE-DUO-1800
  ad0     3823 MB  SMART CF            201011150208AF59AF59 Compact Flash
  ad1    62720 MB  SMART Lite SATA Drive 2010122700A160026002 Disk 1
CB 0             REV 09   710-022597   EF7371           LCC Control Board
....
```

- Run the `show chassis hardware extensive` operational command:

show chassis hardware extensive

```
user@T4000> show chassis hardware extensive
Hardware inventory:
```

```
Item             Version  Part number  Serial number     Description
Chassis                                JN11B3892AHA      T640
Jedec Code:   0x7fb0           EEPROM Version:    0x02
                                S/N:               JN11B3892AHA
Assembly ID:  0x0507           Assembly Version:  00.00
Date:         00-00-0000       Assembly Flags:    0x00
....
```

Verifying the OSS Mapping Feature on SNMP MIBs

Purpose

To verify that the SNMP MIBs are updated with the configured known chassis name.

Action

Run the `show snmp mib` operational commands and verify that the configured known chassis name is displayed in SNMP MIBs when the T4000 chassis is mapped to a T640 chassis:

- Run the `show snmp mib walk system` operational command:

  show snmp mib walk system

```
user@T4000> show snmp mib walk system
sysDescr.0    = Juniper Networks, Inc. t640 internet router, kernel JUNOS 12.3-...Juniper
Networks, Inc.
sysObjectID.0 = jnxProductNameT640
...
```

- Run the `show snmp mib walk jnxBoxAnatomy` operational command:

  show snmp mib walk jnxBoxAnatomy

```
user@T4000> show snmp mib walk jnxBoxAnatomy
jnxBoxClass.0 = jnxProductLineT640.0
jnxBoxDescr.0 = Juniper t640 Internet Backbone Router
jnxBoxSerialNo.0 = JN11B3892AHA
jnxBoxRevision.0
....
```

Meaning

On configuring the OSS mapping feature, the OSS maps the T4000 chassis to a T640 chassis, thereby preventing requalification of the new chassis.

RELATED DOCUMENTATION

oss-map

show chassis oss-map

9

CHAPTER

Configuration Statements and Operational Commands

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- Junos CLI Reference

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- Configuration Statements
- Operational Commands