JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Link and Multilink Services Interfaces User Guide for Routing Devices

Published
2025-12-14

JUNOS

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 Link and Multilink Services Interfaces User Guide for Routing Devices
Copyright © 2025 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 | vi

1 Overview

Understanding Link and Multilink Services | 2

- Link and Multilink Services Overview | 2
- Multilink and Link Services PICs Overview | 3
- Multilink Interfaces on Channelized MICs Overview | 7

2 Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles

Configuring Link and Multilink Services Bundles | 10

- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 10
- Configuring the Number of Bundles on Link Services PICs | 11
- Configuring the Links in a Multilink or Link Services Bundle | 12
- Example: Configuring a Link Services Interface with Two Links | 14
- Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types | 16
- Enabling MLPPP Link Fragmentation and Interleaving | 24
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces | 30

Configuring MLPPP | 34

- Understanding MLPPP | 34
- Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers | 35
- Example: Configuring an MLPPP Bundle on ACX Series | 42
  - Requirements | 42
  - Overview | 42
  - Configuration | 42
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on ACX Series | 44
- Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP | 48

Understanding Multiclass MLPPP | 52

Configuring Multiclass MLPPP on LSQ Interfaces | 53

Examples: Bundling Multiple PPP Links on a Single Link Using MLPPP | 54

Example: Configuring a Multilink Interface with MLPPP | 55

Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces | 56

Example: Configuring an MLPPP Bundle | 58

> Requirements | 59
> Overview | 59
> Configuration | 59
> Verification | 63

Example: Configuring a Link Services Interface with MLPPP | 64

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces | 65

> Requirements | 65
> Overview | 65
> Configuration | 66
> Verification | 75

Examples: Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR | 88

Example: Configuring a Multilink Interface with MLFR FRF.15 | 88

Example: Configuring Multilink Frame Relay FRF.16 | 90

> Requirements | 90
> Overview | 90
> Configuration | 90
> Verification | 95

Example: Configuring Multilink Frame Relay FRF.15 | 96

> Requirements | 96
> Overview | 96
> Configuration | 96
> Verification | 100

Configuring DLCIs on Link Services Logical Interfaces | 100

Example: Configuring a Link Services PIC with MLFR FRF.16 | 102

Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces | **103**

Requirements | **104**

Overview | **104**

Configuration | **105**

Verification | **111**

Example: Configuring Link Interfaces on Channelized MICs | 122

Requirements | **122**

Overview | **122**

Configuration on SDH OC3/STM1 (Multi-Rate) MIC with SFP | **123**

Verification | **128**

Unified ISSU on Inline LSQ Interfaces Overview | 129

3 **Configuring the Physical and Logical Interfaces in a Link and Multilink Configuration**

Configuring Link Services Physical Interfaces | 134

Configuring Link and Multilink Services Logical Interfaces | 134

Multilink and Link Services Logical Interface Configuration Overview | **135**

Configuring Encapsulation for Multilink and Link Services Logical Interfaces | **136**

Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces | **137**

Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces | **139**

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces | **140**

Configuring MRRU on Multilink and Link Services Logical Interfaces | **141**

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces | **142**

Configuring CoS on Link Services Interfaces | **143**

4 **Configuration Statements and Operational Commands**

Junos CLI Reference Overview | 149

About This Guide

Use this guide to configure and monitor multilink bundles. Multilink bundles coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the member links.

1
CHAPTER

Overview

IN THIS CHAPTER

- Understanding Link and Multilink Services | **2**

Understanding Link and Multilink Services

IN THIS SECTION

- Link and Multilink Services Overview | 2
- Multilink and Link Services PICs Overview | 3
- Multilink Interfaces on Channelized MICs Overview | 7

Link and Multilink Services Overview

IN THIS SECTION

- Standards | 3

Multilink-based protocols enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of a multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members. In addition to providing incremental bandwidth, bundling multiple links can add a level of fault tolerance to your dedicated access service, because you can implement bundling across multiple PICs, protecting against the failure of any single PIC.

Junos OS supports several multilink-based protocols including Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR). MLPPP enables you to bundle multiple PPP links into a single logical link. MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single logical link. MLPPP and MLFR provide service option granularity between low-speed T1 and E1 services and higher-speed T3 and E3 services. You use MLPPP and MLFR to increase bandwidth in smaller, more cost-effective increments.

The multiclass extension to the MLPPP extension enables multiple classes of service using MLPPP. For more information, see RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. The Junos OS PPP implementation does not support the negotiation of address field compression and protocol field compression PPP NCP options. The software always sends a full 4-byte PPP header.

Standards

The standards for MLPPP, MLFR FRF.15, and MLFR FRF.16 are defined in the following specifications:

- RFC 1990, *The PPP Multilink Protocol (MP)*
- FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*
- FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

> **NOTE:** Endpoint Discriminator Class compatibility checking is enabled on MLPPP interfaces. Prior to Junos OS Release 8.0, when a Juniper Networks router received an unsupported Endpoint Discriminator Class message from an MLPPP session peer, it returned an ACK response.

Multilink and Link Services PICs Overview

IN THIS SECTION

- Support on Service PICs | 4
- Support on Interface Types | 5

Each Multilink Services or Link Services PIC can support a number of *bundles*. A bundle can contain up to eight individual *links*.

For Multilink Services PICs, the links can be T1, E1, or DS0 physical interfaces, and each link is associated with a logical unit number that you configure. For Link Services PICs, the links can be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1 interfaces, or channelized IQ interfaces. For MLFR FRF.16 bundles, each link is associated with a channel number that you configure.

You must configure a link before it can join a bundle. Each bundle should consist solely of one type of link; the mixing of physical interfaces of differing speeds within a bundle is not supported.

Three versions of Multilink Services and three versions of Link Services PICs are available, as shown in Table 1 on page 4. The PIC hardware is identical, except for different faceplates that enable you to identify which version you are installing. The software limits the unit numbers and maximum number of physical interfaces you assign to the PIC.

Table 1: Multilink and Link Services PIC Capacities

PIC Capacity	Unit Numbers	Maximum Number of T1/DS0 Interfaces	Maximum Number of E1 Interfaces
4-bundle PIC	0 through 3	32 links	32 links
32-bundle PIC	0 through 31	256 links	219 links
128-bundle PIC	0 through 127	292 links	219 links

A single PIC can support an aggregate bandwidth of 450 megabits per second (Mbps).

You can configure a larger number of links, but the Multilink Services and Link Services PICs can reliably process only 450 Mbps of traffic. A higher rate of traffic might degrade performance.

NOTE: In Junos OS releases 9.0 and above you are not allowed to configure a unit number greater than the maximum unit number available on your link services PIC. Attempting to do so will cause an error message.

Support on Service PICs

Junos OS supports multilink-based protocols on services PICs such as the Multilink Services PIC and the Link Services PIC, as well as the link services intelligent queuing (IQ) and voice services configured on the Adaptive Services (AS) and MultiServices PICs. For more information about link services IQ, see *Layer 2 Service Package Capabilities and Interfaces*. For more information about voice services, see *Configuring Services Interfaces for Voice Services*.

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

For more information about MLPPP-based services MICs, see "Multilink Interfaces on Channelized MICs Overview" on page 7.

The Link Services and Multilink Services PICs support the following encapsulation types:

- MLPPP
- MLFR

Starting with Junos OS Release 12.1, support for the following encapsulation types and protocols has been extended to the MX240, MX480, and MX960 routers with Multiservices DPCs:

- MLPPP
- Multiclass MLPPP
- MLFR end-to-end (FRF.15)
- MLFR UNI NNI (FRF.16) (also referred to as MFR)
- Compressed Real-Time Transport Protocol (CRTP)

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

At the logical unit level, the Multilink Services and Link Services PICs support the MLPPP and MLFR Frame Relay Forum (FRF) 15 encapsulation types. At the physical interface level, the Link Services PIC also supports the MLFR FRF.16 encapsulation type.

NOTE: On M Series Multiservice Edge Routers, only one DS3 link is allowed in an MLFR bundle. MLPPP bundles can include two DS3 links.

On ACX Series routers, even if the PIC can support up to 4xDS3 total throughput, each aggregate can only run a volume of traffic equal to one DS3 in bandwidth. Aggregating DS3 links is not supported.

Support on Interface Types

MLPPP and MLFR FRF.15 are supported on interface types `ml-`*`fpc/pic/port`*, `ls-`*`fpc/pic/port`*, and `lsq-`*`fpc/pic/port`*. For MLFR FRF.15, multiple permanent virtual circuits (PVCs) are combined into one aggregated virtual circuit (AVC). This provides fragmentation over multiple PVCs on one end and reassembly of the AVC on the other end.

MLFR FRF.16 is supported on a channelized interface, `ls-`*`fpc/pic/port`*`:`*`channel`*, which denotes a single MLFR FRF.16 bundle. For MLFR FRF.16, multiple links are combined to form one logical link. Packet fragmentation and reassembly occur on a per-VC basis. Each bundle can support multiple VCs. Link Services PICs can support up to 256 DLCIs per MLFR FRF.16 bundle. The physical connections must be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1, or

channelized IQ interfaces. When you bundle channelized interfaces using the link services interface, the channelized interfaces require M Series Enhanced Flexible PIC Concentrators (FPCs).

The `ml-` interface type is used to configure interfaces on the Multilink Services PIC and does not support class-of-service (CoS) features. The `ls-` interface type is used for limited CoS configurations on the Link Services PIC, and the `lsq-` interface type is used for full CoS configurations on the Adaptive Services and MultiServices PICs. The bundle interfaces are configured on the Multiservices DPC as link services IQ (lsq) interfaces and virtual LSQ redundancy (rlsq) interfaces.

For link services IQ (`lsq`) interfaces, Junos OS CoS components are fully supported and are handled normally on M Series and T Series routers, as described in the Junos OS Class of Service User Guide for Routing Devices. For more information on link services IQ configuration, see *Layer 2 Service Package Capabilities and Interfaces*.

When running MLPPP or MLFR on a non-QPP interface, you cannot mix logical units that are members of an aggregate with logical units configured using other families, such as `inet`. For example, the following configuration is not valid:

```
interface e3-0/0/0 {
    encapsulation frame-relay;
    unit 99 {
        dlci 99;
        family mlfr-end-to-end {
            bundle ls-0/0/0.1;
        }
    }
    unit 100 {   ## mixes mlfr with family inet
        dlci 100;
        family inet {
            address 192.168.164.53/30;
        }
    }
}
```

SEE ALSO

Configuring Link and Multilink Services Logical Interfaces | 134

Multilink Interfaces on Channelized MICs Overview

Multiservices Modular Interface Cards (MICs) enable you to perform multiple services on the same MIC by configuring a set of services and applications such as voice services and Layer 2 Tunneling Protocol (L2TP) services. On Juniper Networks MX Series 5G Universal Routing Platforms, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way. The Multilink interfaces are hosted on a channelized MIC. The bundle interfaces are configured on Multiservices DPC as virtual LSQ redundancy (rlsq) interfaces.

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

The following encapsulations, interfaces, protocol, and packet types are supported on the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports *Priority-based Flow Control* (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and Multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Multilink Frame Relay (MFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) non multilink MLPPP and MLFR packets.

Layer 2 services and voice services functionality are implemented on the Multiservices Dense Port Concentrators, which supports the following two kinds of traffic that are routed by the Packet Forwarding Engine:

- Customer-end to provider-end (also, known as customer traffic)—Here, the Multilink fragments from the customer end arrive at the Multiservices interfaces configured on the channelized MIC. These fragments are then transmitted to the Multiservices DPC for Layer 2 processing such as CoS and are reassembled by the Multiservices software running on the Multiservices DPC. These reassembled packets are sent to the Packet Forwarding Engine where they go through the regular router lookup process and are finally sent over the Internet to the provider end. The voice packets also go through the same process.

- Provider-end to customer-end (also, known as Internet traffic)—Here, the data packets that are sent from the Internet provider end are received at any generic ingress interface in the Packet Forwarding Engine. These packets are then sent to the Multiservices DPC for Layer 2 processing. The Multiservices software running on Multiservices DPC fragment these data packets and send it to the Packet Forwarding Engine. These Multilink fragments are sent over the channelized MIC interfaces to the customer end. The voice packets also go through the same process.

NOTE: All the features that are supported on Multilink and Link Services PICs are also supported on the Multilink Services or Link Services MICs. For more information about Multilink and Link Services PICs, see "Multilink and Link Services PICs Overview" on page 3.

Support for the following encapsulations, interfaces, protocol, and packet types are now extended to the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports priority-based flow control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI) and Consortium LMI (C-LMI) for data and control packets.
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) on multilink MLPPP and MLFR packets—Reduces delay and *jitter* on links by breaking up large data packets and interleaving delay-sensitive voice packets with the resulting smaller packets.

RELATED DOCUMENTATION

Configuring MLPPP | **34**

Example: Configuring an MLPPP Bundle | **58**

Example: Configuring Multilink Frame Relay FRF.15 | **96**

Example: Configuring Multilink Frame Relay FRF.16 | **90**

Example: Configuring Link Interfaces on Channelized MICs | **122**

2

CHAPTER

Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles

IN THIS CHAPTER

- Configuring Link and Multilink Services Bundles | **10**
- Configuring MLPPP | **34**
- Examples: Bundling Multiple PPP Links on a Single Link Using MLPPP | **54**
- Examples: Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR | **88**
- Example: Configuring Link Interfaces on Channelized MICs | **122**
- Unified ISSU on Inline LSQ Interfaces Overview | **129**

Configuring Link and Multilink Services Bundles

IN THIS SECTION

- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 10
- Configuring the Number of Bundles on Link Services PICs | 11
- Configuring the Links in a Multilink or Link Services Bundle | 12
- Example: Configuring a Link Services Interface with Two Links | 14
- Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types | 16
- Enabling MLPPP Link Fragmentation and Interleaving | 24
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces | 30

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links

MX240, MX480, and MX960 Universal Routing Platforms support MLPPP and MLFR multilink encapsulations. MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

You configure multilink bundles as logical units or channels on the link services interface `lsq-0/0/0`:

- With MLPPP and MLFR FRF.15, multilink bundles are configured as logical units on `lsq-0/0/0`—for example, `lsq-0/0/0.0` and `lsq-0/0/0.1`.
- With MLFR FRF.16, multilink bundles are configured as channels on `lsq-0/0/0`—for example, `lsq-0/0/0:0` and `lsq-0/0/0:1`.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. Depending on the system license and hardware, you can create up to 1023 multilink bundles per Chassis, and on each multilink bundle add up to 8 constituent links. See "Multilink and Link Services PICs Overview" on page 3 for more information.

The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.
- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle, and only interfaces with a Frame Relay encapsulation can be added to an MLFR bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

Configuring a multilink bundle on the two serial links increases the bandwidth by 70 percent from approximately 1 Mbps to 1.7 Mbps and prepends each packet with a multilink header as specified in the FRF.12 standard. To increase the bandwidth further, you can add up to eight serial links to the bundle. In addition to a higher bandwidth, configuring the multilink bundle provides load balancing and redundancy. If one of the serial links fails, traffic continues to be transmitted on the other links without any interruption. In contrast, independent links require routing policies for load balancing and redundancy. Independent links also require IP addresses for each link as opposed to one IP address for the bundle. In the routing table, the multilink bundle is represented as a single interface.

Starting with Junos OS Release 13.3, if you attempt to delete or deactivate a static inline service (`si`) MLPPP bundle interface that is still referenced by a member link interface, which could be PPPoE (`pp0`) or `si` logical interfaces, and commit the configuration, the commit operation fails. You must reactivate such MLPPP bundle interface before committing the settings. Alternatively, you must ensure that member links do not refer a static MLPPP bundle before you delete or deactivate the bundle. This method of deactivation and reactivation of an MLPPP bundle is not applicable for interfaces other than `si-` interfaces, such as link services IQ (`lsq-`) and virtual LSQ redundancy (`rlsq-`) interfaces.

SEE ALSO

Example: Configuring Link Interfaces on Channelized MICs | **122**

Example: Configuring an MLPPP Bundle | **58**

Example: Configuring Multilink Frame Relay FRF.15 | **96**

Example: Configuring Multilink Frame Relay FRF.16 | **90**

Configuring the Number of Bundles on Link Services PICs

You can combine MLFR FRF.16, MLPPP, and MLFR FRF.15 bundles on a single Link Services PIC. For a sample configuration, see "Example: Configuring a Link Services Interface with Two Links" on page 14.

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

```
mlfr-uni-nni-bundles number;
```

Each Link Services PIC can accommodate a maximum of 256 MLFR UNI NNI bundles. For more information, see the Junos OS Administration Library for Routing Devices.

A link can associate with one link services bundle only. All Link Services PICs support up to 256 single-link bundles and up to 256 DLCIs. For an example configuration, see the configuration examples.

NOTE: When one or more links in a bundle are put in loopback, reassembly buffering and hence processing are reduced so as to not affect other bundles. This prevents packet loss on other bundles, while reducing the reassembly buffers available for the bundle with looped links.

SEE ALSO

Example: Configuring a Link Services Interface with MLPPP | 64

Example: Configuring a Link Services Interface with MLFR FRF.15

Example: Configuring a Link Services PIC with MLFR FRF.16 | 102

Configuring the Links in a Multilink or Link Services Bundle

To complete a multilink or link services interface configuration, you need to configure both the physical interface and the multilink or link services bundle. For multilink interfaces, you configure the link bundle on the logical unit. For link services interfaces, you configure the link bundle as a channel (see Figure 1 on page 13). The physical interface is usually connected to networks capable of supporting MLPPP or MLFR (FRF.15 or FRF.16).

Figure 1: Multilink Interface Configuration

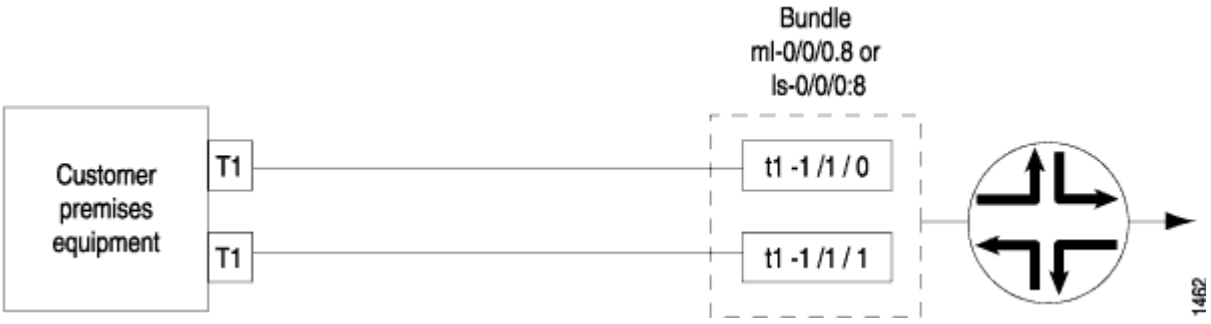

The following sample configuration refers to the topology in Figure 1 on page 13 and configures a multilink or link services bundle over a T1 connection (for which the T1 physical interface is already configured).

1. To configure a physical T1 link for MLPPP, include the following statements at the `[edit interfaces t1-`*fpc*`/`*pic*`/`*port*`]` hierarchy level:

```
unit 0 {
    family mlppp {
        bundle (ml-fpc/pic/port | ls-fpc/pic/port);
    }
}
```

   You do not need to configure an IP address on this link.

   To configure a physical T1 link for MLFR FRF.16, include the following statements at the `[edit interfaces t1-`*fpc*`/`*pic*`/`*port*`]` hierarchy level:

```
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle ls-fpc/pic/port:channel;
    }
}
```

   You do not need to configure an IP address or a DLCI on this link.

2. To configure the logical address for the MLPPP, MLFR FRF.15, or MLFR FRF.16 bundle, include the `address` and `destination` statements:

```
address address {
    destination address;
}
```

You can include these statements at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family inet]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet]`

When you add statements such as `mrru` to the configuration and commit, the T1 interface becomes part of the multilink bundle.

NOTE: For MLPPP and MLFR (FRF.15 and FRF.16) links, you must specify the subnet address as `/32` or `/30`. Any other subnet designation is treated as a mismatch.

SEE ALSO

Link and Multilink Services Overview | 2

Example: Configuring an MLPPP Bundle | 58

Example: Configuring a Link Services Interface with Two Links

This example uses the MLFR UNI NNI protocol between Router A and Router B and logically connects link services bundles `ls-1/1/0.3` and `ls-0/0/0.10`, as specified in Table 2 on page 14.

Table 2: Link Services Bundle

Router A	Router B
t1-0/1/0 (ls-1/1/0:3)	t1-0/3/0 (ls-0/0/0:10)
t1-0/1/1 (ls-1/1/0:3)	t1-0/3/1 (ls-0/0/0:10)

For LMI to work properly, you must configure one router to be a DCE.

Configuration on Router A

```
[edit interfaces]
ls-1/1/0:3 {
    dce;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        dlci 16;
        family inet {
            address 10.3.3.1/32 {
                destination 10.3.3.2;
            }
        }
    }
}
t1-0/1/0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/1/0:3;
        }
    }
}
t1-0/1/1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/1/0:3;
        }
    }
}
```

Configuration on Router B

```
[edit interfaces]
ls-0/0/0:10 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        dlci 16;
        family inet {
```

```
            address 10.3.3.2/32 {
                destination 10.3.3.1;
            }
        }
    }
}
t1-0/3/0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-0/0/0:10;
        }
    }
}
t1-0/3/1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-0/0/0:10;
        }
    }
}
```

SEE ALSO

> *encapsulation (Physical Interface)*
>
> Configuring Link Services Physical Interfaces | 134

Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types

```
[edit chassis]
fpc 1 {
    pic 3 {
        mlfr-uni-nni-bundles 4;
    }
}
[edit interfaces]
```

```
t1-0/2/0:0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/3/0:0;
        }
    }
}
t1-0/2/0:1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/3/0:0;
        }
    }
}
t1-0/2/0:5 {
    unit 0 {
        family mlppp {
            bundle ls-1/3/0.2;
        }
    }
}
t1-0/2/0:6 {
    unit 0 {
        family mlppp {
            bundle ls-1/3/0.2;
        }
    }
}
t1-0/2/0:7 {
    encapsulation frame-relay;
    unit 0 {
        dlci 20;
        family mlfr-end-to-end {
            bundle ls-1/3/0.1;
        }
    }
}
t1-0/2/0:8 {
    encapsulation frame-relay;
    unit 0 {
        dlci 20;
```

```
        family mlfr-end-to-end {
            bundle ls-1/3/0.1;
        }
    }
}
t1-0/2/0:10 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
t3-1/0/0 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.2;
        }
    }
}
lsq-1/1/0 {
    unit 0 {
        encapsulation multilink-ppp;
        compression {
            rtp {
                f-max-period 100;
                queues [ q1 q2 ];
                port minimum 2000 maximum 6000;
            }
        }
        family inet {
            address 10.5.5.5/24;
        }
    }
    unit 1 {
        encapsulation multilink-ppp;
        compression {
            rtp {
                port minimum 2000 maximum 6000;
            }
```

```
        }
        family inet {
            address 10.6.6.1/24;
        }
    }
    unit 2 {
        encapsulation multilink-ppp;
        compression {
            rtp {
                port minimum 2000 maximum 6000;
            }
        }
        family inet {
            address 10.9.9.1/24;
        }
    }
}
t1-1/2/0 {
    no-keepalives;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.1;
        }
    }
}
ls-1/3/0 {
    unit 1 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.1.4.1/24;
        }
    }
    unit 2 {
        encapsulation multilink-ppp;
        family inet {
            address 10.7.4.1/24;
        }
    }
}
ls-1/3/0:0 {
    encapsulation multilink-frame-relay-uni-nni;
    mlfr-uni-nni-bundle-options {
        debug-flags 15;
```

```
    }
    unit 0 {
        dlci 20;
        family inet {
            address 10.5.4.1/24;
        }
    }
}
[edit routing-options]
static {
    route 10.12.12.0/24 next-hop 10.1.1.9;
}
```

On Router B:

```
[edit chassis]
fpc 1 {
    pic 3 {
        mlfr-uni-nni-bundles 4;
    }
}
[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.1.1/24;
        }
    }
}
so-0/1/1 {
    encapsulation ppp;
    unit 0 {
        family inet {
            address 10.7.7.7/24;
        }
    }
}
t1-0/2/0:0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/3/0:0;
```

```
        }
    }
}
t1-0/2/0:1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/3/0:0;
        }
    }
}
t1-0/2/0:5 {
    no-keepalives;
    unit 0 {
        family mlppp {
            bundle ls-1/3/0.2;
        }
    }
}
t1-0/2/0:6 {
    no-keepalives;
    unit 0 {
        family mlppp {
            bundle ls-1/3/0.2;
        }
    }
}
t1-0/2/0:7 {
    dce;
    encapsulation frame-relay;
    unit 0 {
        dlci 20;
        family mlfr-end-to-end {
            bundle ls-1/3/0.1;
        }
    }
}
t1-0/2/0:8 {
    dce;
    encapsulation frame-relay;
    unit 0 {
        dlci 20;
        family mlfr-end-to-end {
```

```
            bundle ls-1/3/0.1;
        }
    }
}
t1-0/2/0:10 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
t3-0/3/0 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.2;
        }
    }
}
ge-1/0/0 {
    unit 0 {
        family inet {
            address 10.2.2.1/24;
        }
    }
}
lsq-1/1/0 {
    unit 0 {
        compression {
            rtp {
                port minimum 2000 maximum 6000;
            }
        }
        family inet {
            address 10.5.5.1/24;
        }
    }
    unit 1 {
        encapsulation multilink-ppp;
        compression {
```

```
            rtp {
                port minimum 16384 maximum 20102;
            }
        }
        family inet {
            address 10.3.4.1/24;
        }
    }
    unit 2 {
        encapsulation multilink-ppp;
        compression {
            rtp {
                port minimum 2000 maximum 6000;
            }
        }
        family inet {
            address 10.9.9.9/24;
        }
    }
}
t1-1/2/2 {
    no-keepalives;
    unit 0 {
        family mlppp {
            bundle ls-1/3/0.1;
        }
    }
}
t1-1/2/3 {
    no-keepalives;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.1;
        }
    }
}
ls-1/3/0 {
    unit 1 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.1.4.4/24;
        }
        family iso;
```

```
    }
    unit 2 {
        encapsulation multilink-ppp;
        family inet {
            address 10.7.4.4/24;
        }
    }
}
ls-1/3/0:0 {
    dce;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        dlci 20;
        family inet {
            address 10.5.4.4/24;
        }
    }
}
[edit routing-options]
static {
    route 10.12.12.0/24 next-hop 10.3.4.4;
}
```

SEE ALSO

> Configuring Link Services Physical Interfaces | **134**
>
> *encapsulation (Physical Interface)*

Enabling MLPPP Link Fragmentation and Interleaving

MLPPP enables you to bundle multiple PPP links into a single multilink bundle. MLPPP bundle support on an inline LSQ interface is identical to a non-inline LSQ interface, because the configuration to enable fragmentation, link fragmentation and interleaving (LFI), and timeout is identical.

Priority scheduling on a multilink bundle determines the order in which an output interface transmits traffic from an output queue. The queues are serviced in a weighted round-robin fashion. But when a queue containing large packets starts using the multilink bundle, small and delay-sensitive packets must wait their turn for transmission. Because of this delay, some slow links, such as T1 and E1, can become useless for delay-sensitive traffic.

Link fragmentation and interleaving (LFI) solves this problem. It reduces delay and jitter on links by fragmenting large packets and interleaving delay-sensitive packets with the resulting smaller packets for simultaneous transmission across multiple links of a multilink bundle.

To configure schedule maps and fragmentation maps for MLPPP LFI:

1. Assign each forwarding class to an internal queue number by including the `forwarding-classes` statement at the `[edit class-of-service]` hierarchy level.

```
[edit class-of-service]
forwarding-classes {
    queue queue-number  class-name;
}
```

For example, to set four output transmission queues:

```
[edit class-of-service]
user@host# set forwarding-classes queue 0 be
user@host# set forwarding-classes queue 1 ef
user@host# set forwarding-classes queue 2 af
user@host# set forwarding-classes queue 3 nc
```

2. To set a per-forwarding class fragmentation threshold, include the `fragment-threshold` statement in the `fragmentation-maps`.

```
[edit class-of-service]
fragmentation-maps {
    map-name{
        forwarding-class class-name {
            fragment-threshold bytes;
        }
    }
}
```

For example, to create two fragmentation maps and set a per-forwarding class fragmentation threshold:

```
[edit class-of-service]
user@host# set fragmentation-maps fragmap-1 forwarding-class af fragment-threshold 320
user@host# set fragmentation-maps fragmap-1 forwarding-class be fragment-threshold 256
user@host# set fragmentation-maps fragmap-1 forwarding-class ef fragment-threshold no-
```

```
fragmentation
user@host# set fragmentation-maps fragmap-2 forwarding-class af fragment-threshold 192
user@host# set fragmentation-maps fragmap-2 forwarding-class be fragment-threshold 320
user@host# set fragmentation-maps fragmap-2 forwarding-class ef fragment-threshold 192
user@host# set fragmentation-maps fragmap-2 forwarding-class nc fragment-threshold no-
fragmentation
```

The `fragment-threshold` statement in the LSQ bundle logical interface configuration applies to all forwarding classes. The `fragment-threshold` statement in `fragmentation-maps` for a particular `forwarding class`, if present, overrides the statement configured in the LSQ bundle logical interface for that class. If `fragment-threshold` is not configured anywhere in the configuration, packets are still fragmented if `fragment-threshold` exceeds the smallest maximum transmission unit (*MTU*) or maximum received reconstructed unit (*MRRU*) of all links in the bundle.

3. Configure transmission scheduling parameters.

```
[edit class-of-service scheduler scheduler-name]
schedulers {
    scheduler-name {
        priority priority-level;
        transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
    }
}
```

For example, to set the transmit-rate percentage and the priority-level:

```
[edit class-of-service scheduler af-scheduler]
user@host# set transmit-rate percent 30
user@host# set priority low
```

```
[edit class-of-service scheduler be-scheduler]
user@host# set transmit-rate percent 20
user@host# set priority low
```

```
[edit class-of-service scheduler ef-scheduler]
user@host# set transmit-rate percent 35
user@host# set priority strict-high
```

```
[edit class-of-service scheduler nc-scheduler]
user@host# set transmit-rate percent 15
user@host# set priority high
```

4. After defining a scheduler, associate it with a specified forwarding class by including it in a *scheduler-map*.

```
[edit class-of-service]
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
```

For example, to associate the af-scheduler, be-scheduler, ef-scheduler, and nc-scheduler schedulers, with the af, be, ef, and nc forwarding-classes:

```
[edit class-of-service]
user@host# set scheduler-maps sched-map1 forwarding-class af scheduler af-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class be scheduler be-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class ef scheduler ef-scheduler
user@host# set scheduler-maps sched-map1 forwarding-class nc scheduler nc-scheduler
```

5. Configure traffic shaping and scheduling profiles.

```
[edit class-of-service]
traffic-control-profiles {
    profile-name {
        guaranteed-rate (percent percentage | rate) <burst-size bytes>;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate) <burst-size bytes>;
    }
}
```

For example, to set the traffic-control policies:

```
[edit class-of-service traffic-control-policies m1-tcp1]
user@host# set guaranteed-rate 1m
user@host# set scheduler-map sched-map1
user@host# set shaping-rate 1m
```

6. Configure interface-specific CoS properties for incoming packets.

```
[edit class-of-service]
interfaces {
    interface-name {
        unit logical-unit-number {
            fragmentation-map map-name;
            output-traffic-control-profile profile-name;
        }
    }
}
```

For example, to apply the specified CoS traffic control profile (traffic scheduling and shaping configuration objects) to the output traffic at the logical interface:

```
[edit class-of-service]
user@host# set interfaces lsq-0/1/0 unit 100 fragmentation-map fragmap-1 output-traffic-
control-profile m1-tcp1
```

The following partial configuration shows when the fragment threshold for low priority queues inherits from the fragment threshold configured in the bundle IFL and will have the value of 640.

```
[edit class-of-service]
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
fragmentation-maps {
    fragmap-3 {
        forwarding-class ef {
            no-fragmentation;
        }
    }
}
schedulers {
    af-scheduler {
        transmit-rate percent 30;
        priority low;
    }
    be-scheduler {
        transmit-rate percent 20;
        priority low;
    }
    ef-scheduler {
        transmit-rate percent 35 rate-limit;
        priority strict-high;
    }
    nc-scheduler {
        transmit-rate percent 15;
        priority high;
    }
    }
....
```

SEE ALSO

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 10

Inline MLPPP for WAN Interfaces Overview

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces | **65**

Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces | **103**

Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces

IN THIS SECTION

- Configuring LFI with DLCI Scheduling | **31**

For link services FRF.15 and MLPPP interfaces only, you can configure link fragment interleaving (LFI). LFI reduces excessive delays of Frame Relay packets by fragmenting long packets into smaller packets and interleaving them with real-time frames. This allows real-time and non-real-time data frames to be carried together on lower-speed links without causing excessive delays to the real-time traffic. When the peer interface receives the smaller fragments, it reassembles the fragments into their original packet. For example, short delay-sensitive packets, such as packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.

NOTE: All Link Services PICs (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle) support up to 256 link services interfaces with LFI enabled, if those link services interfaces contain only one constituent link each. For the Link Services PIC, multiple-link LFI bundles are simply multilink bundles, and are limited based on the type of PIC (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle).

In addition, the multilink bundles you configure subtract from the total of 256 possible LFI-enabled link services interfaces. For example, if a 32-multilink bundle Link Services PIC has 24 multilink bundles configured and active, then you can configure 256 – 24 = 232 LFI-enabled link services interfaces, each with a single constituent link.

For link services IQ interfaces (`lsq`), the `interleave-fragments` statement is not valid. Instead, you can enable LFI by configuring fragmentation maps. For more information, see *Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces*.

You can configure multiple links in a bundle and configure packet interleaving. However, if you use packet interleaving, high-priority, nonmultilink-encapsulated packets use a hash-based algorithm to choose a single link.

Per-bundle CoS queuing is supported on link services IQ interfaces (lsq). For more information about link services IQ interfaces, see *Layer 2 Service Package Capabilities and Interfaces*.

The Junos OS supports end-to-end fragmentation in compliance with the FRF.12 *Frame Relay Fragmentation Implementation Agreement* standard. Unlike user-to-network interface (UNI) and network-to-network (NNI) fragmentation, end-to-end supports fragmentation only at the endpoints.

By default, packet interleaving is disabled. To enable packet interleaving, include the interleave-fragments statement:

```
interleave-fragments;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring LFI with DLCI Scheduling

For Link Services and Channelized DS3 IQ PICs, you can configure LFI and DLCI scheduling. For channelized DS3 interfaces, LFI is supported with FRF.15 only, and on M10i and M20 platforms only.

Configuring LFI with DLCI scheduling enables packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

For more information about associating a scheduler with a DLCI, see the Junos OS Class of Service User Guide for Routing Devices.

Example: Configuring LFI with DLCI Scheduling

Configure packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

```
[edit interfaces]
ls-1/0/0 {
    unit 1 {
```

```
        encapsulation multilink-frame-relay-end-to-end;
        interleave-fragments;
        family inet {
            address 192.168.5.2/32 {
                destination 192.168.5.3;
            }
        }
    }
    t3-1/0/0:1 {
        per-unit-scheduler;
        unit 0 {
            dlci 16;
            encapsulation multilink-frame-relay-end-to-end;
            family mlfr-end-to-end {
                bundle ls-1/0/0.1;
            }
        }
    }
    [edit class-of-service]
    interfaces {
        t3-1/0/0:1 {
            unit 0 {
                scheduler-map sched-map-logical-0;
                shaping-rate 10m;
            }
            unit 1 {
                scheduler-map sched-map-logical-1;
                shaping-rate 20m;
            }
        }
    }
    scheduler-maps {
        sched-map-logical-0 {
            forwarding-class best-effort scheduler sched-best-effort-0;
            forwarding-class assured-forwarding scheduler sched-bronze-0;
            forwarding-class expedited-forwarding scheduler sched-silver-0;
            forwarding-class network-control scheduler sched-gold-0;
        }
        sched-map-logical-1 {
            forwarding-class best-effort scheduler sched-best-effort-1;
            forwarding-class assured-forwarding scheduler sched-bronze-1;
            forwarding-class expedited-forwarding scheduler sched-silver-1;
            forwarding-class network-control scheduler sched-gold-1;
```

```
        }
        schedulers {
            sched-best-effort-0 {
                transmit-rate 4m;
            }
            sched-bronze-0 {
                transmit-rate 3m;
            }
            sched-silver-0 {
                transmit-rate 2m;
            }
            sched-gold-0 {
                transmit-rate 1m;
            }
            sched-best-effort-1 {
                transmit-rate 8m;
            }
            sched-bronze-1 {
                transmit-rate 6m;
            }
            sched-silver-1 {
                transmit-rate 4m;
            }
            sched-gold-1 {
                transmit-rate 2m;
            }
        }
    }
}
```

SEE ALSO

- Link and Multilink Services Overview | 2
- Configuring Link and Multilink Services Bundles | 10
- Configuring DLCIs on Link Services Logical Interfaces | 100

Configuring MLPPP

IN THIS SECTION

- Understanding MLPPP | 34
- Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers | 35
- Example: Configuring an MLPPP Bundle on ACX Series | 42
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on ACX Series | 44
- Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP | 48
- Understanding Multiclass MLPPP | 52
- Configuring Multiclass MLPPP on LSQ Interfaces | 53

Understanding MLPPP

IN THIS SECTION

- MLPPP Support on ACX Series Routers | 35

Multilink Point-to-Point Protocol (MLPPP) enables you to bundle multiple PPP links into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1 and E1 links.

You configure multilink bundles as logical units or channels on the link services interface. With MLPPP, multilink bundles are configured as logical units on the link service interface—for example, `lsq-0/0/0.0`, `lsq-0/0/0.1`. After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated.

The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.
- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle.

- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

With MLPPP bundles, you can use PPP Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP) for secure transmission over the PPP interfaces. For more information, see *Configuring the PPP Challenge Handshake Authentication Protocol* and *Configuring the PPP Password Authentication Protocol*.

MLPPP Support on ACX Series Routers

ACX Series routers support MLPPP encapsulations. MLPPP is supported on ACX1000, ACX2000, ACX2100 routers, and with Channelized OC3/STM1 (Multi-Rate) MICs with SFP and 16-port Channelized E1/T1 Circuit Emulation MIC on ACX4000 routers.

The following table shows the maximum number of multilink bundles you can create on ACX Series routers:

Table 3: Multilink Bundles Supported by ACX Series Routers

ACX Platform	Maximum Bundles	Maximum Links	Maximum Links Per Bundle
ACX2000 ACX2100	16	16	16
ACX4000 ACX-MIC-16CHE1-T1-CE	16	16	16
ACX4000 ACX-MIC-4COC3-1COC12CE	50	336	16
ACX1000	8	8	8

Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers

You can configure MLPPP bundle interfaces with T1/E1 member links. The traffic that is transmitted over the MLPPP bundle interface is spread over the member links in a round-robin manner. If the packet size is higher than the fragmentation size configured on the MLPPP interface, the packet are

fragmented. The fragments are also sent over member links in a round-robin pattern. The PPP control packets received on the interface are terminated on the router. The fragmentation size is configured at the MLPPP bundle-level. This fragmentation size is applied to all the packets on the bundle, regardless of the multilink class.

Multiclass MLPPP segregates the multilink protocol packets in to multiple classes. ACX routers support up to a maximum of four classes. One queue is associated with each of the four classes of multiclass MLPPP. The packets can be classified to be part of one of the classes. These packets take the queue associated with the class. The packets inside a queue are served in first-in first-out (FIFO) sequence.

Multiclass MLPPP is required to provide preferential treatment to high-priority, delay-sensitive traffic. The delay-sensitive smaller real-time frames are classified such that they end up in higher priority queue. While a lower priority packet is being fragmented, if a higher priority packet is enqueued, the lower priority fragmentation is suspended, the higher priority packet is fragmented and enqueued for transmission, and then the lower priority packet fragmentation is resumed.

Traditional LSQ interfaces (anchored on PICs) are supported to combine T1/E1 interfaces in an MLPPP bundle interface. Inline services (si-) interfaces and inline LSQ interfaces are not supported in MLPPP bundles. On ACX routers, MLPPP bundling is performed on the TDM MICs and traditional LSQ model is most effective mechanism. You can configure channelized OC interfaces (`t1-`*`x/y/z:n:m`*, `e1-`*`x/y/z:n`*) as members of an MLPPP bundle interface. A maximum of 16 member links per bundle is supported. The MPLS, ISO, and inet address families are supported. The ISO address family is supported only for IS-IS. You can configure MLPPP bundles on network-to-network interface (NNI) direction of an Ethernet pseudowire. Interleaving using multiclass MLPPP is supported.

Keep the following points in mind when you configure MLPPP bundles on ACX routers:

- The physical links must be of the same type and bandwidth.
- Round-robin packet distribution is performed over the member links.
- To add a T1 or E1 member link to the MLPPP bundle as link services LSQ interfaces, include the `bundle` statement at the `[edit interfaces t1-`*`fpc/pic/port`* `unit` *`logical-unit-number`* `family mlppp]` hierarchy level:

  ```
  [edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
  bundle lsq-fpc/pic/port.logical-unit-number;
  ```

- To configure the link services LSQ interface properties, include the following statements at the `[edit interfaces lsq-`*`fpc/pic/port`* `unit` *`logical-unit-number`*`]` hierarchy level:

  ```
  [edit interfaces lsq-fpc/pic/port unit logical-unit-number]
  encapsulation multilink-ppp;
  fragment-threshold bytes;
  ```

```
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

You can configure the address family as MPLS for the LSQ interfaces in an MLPPP bundle.

- PPP control protocol support depends on the processing of the PPP application for MLPPP bundle interfaces IPv4, Internet Protocol Control Protocol (IPCP), PPP Challenge Handshake Authentication Protocol (CHAP), and Password Authentication Protocol (PAP) applications are supported for PPP.
- Drop timeout configuration is not applicable to ACX routers
- The member links across MICs cannot be bundled. Only physical interfaces on the same MIC can be bundled.
- Fractional T1 and E1 interfaces are not supported. CoS is supported only for full T1 and E1 interfaces. Selective time slots of T1/E1 cannot be used and full T1/E1 interfaces must be used.
- Detailed statistics displayed depend on the parameters supported by the hardware. The counters that are supported by the hardware are displayed with appropriate values in the output of the `show interfaces lsq-fpc/pic/port detail` command.

  In the following sample output, the fields that are displayed with a value of 0 denote the fields that are not supported for computation by ACX routers. In the lsq- interface statistics, non-fragment statistics of the bundle are not accounted. Non-fragments are typically treated as single-fragment frames and counted in the fragment statistics.

```
user@host# show interfaces lsq-1/1/0 detail
Physical interface: lsq-1/1/0, Enabled, Physical link is Up
  Interface index: 162, SNMP ifIndex: 550, Generation: 165
  Description: LSQ-interface
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x0
  Last flapped   : 2015-06-22 19:01:47 PDT (2d 04:56 ago)
  Statistics last cleared: 2015-06-23 05:01:49 PDT (1d 18:56 ago)
  Traffic statistics:
   Input  bytes  :               108824               208896 bps
   Output bytes  :                90185               174080 bps
   Input  packets:                 1075                  256 pps
   Output packets:                 1061                  256 pps
```

```
 IPv6 transit statistics:
  Input  bytes  :                  0
  Output bytes  :                  0
  Input  packets:                  0
  Output packets:                  0
Frame exceptions:
  Oversized frames                 0
  Errored input frames             0
  Input on disabled link/bundle    0
  Output for disabled link/bundle  0
  Queuing drops                    0
Buffering exceptions:
  Packet data buffer overflow      0
  Fragment data buffer overflow    0
Assembly exceptions:
  Fragment timeout                 0
  Missing sequence number          0
  Out-of-order sequence number     0
  Out-of-range sequence number     0
Hardware errors (sticky):
  Data memory error                0
  Control memory error             0

Logical interface lsq-1/1/0.0 (Index 326) (SNMP ifIndex 599) (Generation 177)
  Flags: Up Point-To-Point SNMP-Traps 0x0 Encapsulation: Multilink-PPP
  Last flapped: 2015-06-24 23:57:34 PDT (00:00:51 ago)
  Bandwidth: 6144kbps
  Bundle links information:
    Active bundle links       4
    Removed bundle links      0
    Disabled bundle links     0
  Bundle options:
    MRRU                                2000
    Remote MRRU                         2000
    Drop timer period                   0
    Inner PPP Protocol field compression enabled
    Sequence number format              short (12 bits)
    Fragmentation threshold             450
    Links needed to sustain bundle      3
    Multilink classes                   4
    Link layer overhead                 4.0 %
  Bundle status:
    Received sequence number        0x0
```

```
    Transmit sequence number        0x0
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :          1076        256        484200       921600
      Output:          1061        256        477450       921600
    Network:
      Input :          2182        256        201812       208896
      Output:          2168        256        192029       174080
  IPV6 Transit Statistics         Packets        Bytes
    Network:
      Input :                     0              0
      Output:                     0              0
  Multilink class 0:
    Multilink:
      Input :          1075        256        483750       921600
      Output:          1061        256        477450       921600
    Network:
      Input :          1061        256        477450       921600
      Output:          1075        256        483750       921600
  Multilink class 1:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink class 2:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
```

```
Multilink class 3:
  Multilink:
    Input :             0          0             0            0
    Output:             0          0             0            0
  Network:
    Input :             0          0             0            0
    Output:             0          0             0            0
Link:
  t1-1/1/1.0
    Up time: 00:00:51
    Input :           280         64        126000       230400
    Output:           266         64        119700       230400
  t1-1/1/2.0
    Up time: 00:00:51
    Input :           266         64        119700       230400
    Output:           265         64        119250       230400
  t1-1/1/3.0
    Up time: 00:00:51
    Input :           265         64        119250       230400
    Output:           265         64        119250       230400
  t1-1/1/4.0
    Up time: 00:00:51
    Input :           265         64        119250       230400
    Output:           265         64        119250       230400
Multilink detail statistics:
Bundle:
  Fragments:
    Input :          1076        256        484200       921600
    Output:          1061        256        477450       921600
  Non-fragments:
    Input :             0          0             0            0
    Output:             0          0             0            0
  LFI:
    Input :             0          0             0            0
    Output:             0          0             0            0
Multilink class 0:
  Fragments:
    Input :          1076        256        484200       921600
    Output:          1061        256        477450       921600
  Non-fragments:
    Input :             0          0             0            0
    Output:             0          0             0            0
Multilink class 1:
```

```
      Fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
      Non-fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
    Multilink class 2:
      Fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
      Non-fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
    Multilink class 3:
      Fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
      Non-fragments:
        Input :             0          0             0           0
        Output:             0          0             0           0
  NCP state: inet: Opened, inet6: Not-configured, iso: Opened, mpls: Opened
    Protocol inet, MTU: 1500, Generation: 232, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 9.1.9/24, Local: 9.1.9.18, Broadcast: Unspecified, Generation: 212
    Protocol iso, MTU: 1500, Generation: 233, Route table: 0
      Flags: Is-Primary
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 234, Route table: 0
      Flags: Is-Primary
```

- For modifying the frame checksum (FCS) in the set of T1 options or E1 options on a MLPPP bundle member link, you must remove the member link out of the bundle by deactivating the link or unconfiguring it as a bundle member, and add the link back to the bundle after FCS modification. You must first remove the link from the bundle and modify FCS. If you are configuring FCS for the first time on the member link, specify the value before it is added to the bundle.

The following MLPPP functionalities are not supported on ACX Series routers:

- Member links across MICs.
- Fragmentation per class (only configurable at bundle level).

- IPv6 address family header compression (no address and control field compression [ACFC] or protocol field compression [PFC]).
- Prefix elision as defined in *RFC 2686, The Multi-Class Extension to Multi-Link PPP*.
- A functionality that resembles link fragmentation and interleaving (LFI) can be achieved using multiclass MLPPP (RFC 2686), which interleaves the high priority packets between lower priority packets. This methodology ensures that the delay desitive packets are sent as soon as they arrive. While LFI-classified packets are sent to a specific member link as PPP packets, the ACX implementation of interleaving contains multilink PPP (also referred to as PPP Multilink, MLP, and MP) headers and fragments that are sent on all member links in a round-robin manner.
- PPP over MLPPP bundle interfaces.

Example: Configuring an MLPPP Bundle on ACX Series

IN THIS SECTION

- Requirements | 42
- Overview | 42
- Configuration | 42

Requirements

You require ACX Series routers to configure the following example

Overview

The following is a sample for configuring an MLPPP bundle on ACX Series routers:

Configuration

IN THIS SECTION

- CLI Quick Configuration | 43
- Procedure | 44

CLI Quick Configuration

```
[edit]
user@host# show interfaces
lsq-1/1/0 {
    description LSQ-interface;
        per-unit-scheduler;
        unit 0 {
        encapsulation multilink-ppp;
            mrru 2000;
            short-sequence;
            fragment-threshold 450;
            minimum-links 3;
            multilink-max-classes 4;
            family inet {
            address 9.1.9.18/24
            }
            family iso;
        family mpls;
    }
}
ct1-1/1/1 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/1 {
    encapsulation ppp;
        unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
ct1-1/1/2 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/2 {
    encapsulation ppp;
        unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
```

```
        }
    }
}
ct1-1/1/3 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/3 {
    encapsulation ppp;
        unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
ct1-1/1/4 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/4 {
    encapsulation ppp;
        unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
```

Procedure

Step-by-Step Procedure

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on ACX Series

LSQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure an *N*xT1 bundle using MLPPP, you aggregate *N* different T1 links into a bundle. The *N*xT1 bundle is called a logical interface, because it can represent, for example, a routing adjacency. To aggregate T1 links into a an MLPPP bundle, include the `bundle` statement at the `[edit interfaces t1-`*`fpc/pic/port`* `unit` *`logical-unit-number`* `family mlppp]` hierarchy level:

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```

To configure the LSQ interface properties, include the following statements at the `[edit interfaces lsq-`*`fpc/pic/port`* `unit` *`logical-unit-number`*`]` hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

NOTE: ACX Series routers do not support drop-timeout and link-layer-overhead properties.

The logical link services IQ interface represents the MLPPP bundle. For the MLPPP bundle, there are four associated queues on M Series routers and eight associated queues on M320 and T Series routers. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

For MLPPP, assign a single scheduler map to the link services IQ interface (`lsq`) and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (`lsq`) and to each constituent link..

NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

If the bundle has more than one link, you must include the `per-unit-scheduler` statement at the `[edit interfaces lsq-`*fpc*`/`*pic*`/`*port*`]` hierarchy level:

```
[edit interfaces lsq-fpc/pic/port]
per-unit-scheduler;
```

To configure and apply the scheduling policy, include the following statements at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
interfaces {
    t1-fpc/pic/port unit logical-unit-number {
        scheduler-map map-name;
    }
}
forwarding-classes {
    queue queue-number class-name;
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder | temporal microseconds);
        priority priority-level;
        transmit-rate (rate | remainder) <exact>;
    }
}
```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the Junos OS Class of Service User Guide for Routing Devices.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the `fragmentation-maps` statement at the `[edit class-of-service]` hierarchy level:

```
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            multilink-class number;
        }
    }
}
```

For *N*xT1 bundles using MLPPP, the byte-wise load balancing used in multilink-encapsulated queues is superior to the flow-wise load balancing used in nonencapsulated queues. All other considerations are equal. Therefore, we recommend that you configure all queues to be multilink encapsulated. You do this by including the `fragment-threshold` statement in the configuration. You use the `multilink-class` statement to map a forwarding class into a multiclass MLPPP. For more information about fragmentation maps, see *Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces*.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on one of the *N* different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the `[edit class-of-service fragmentation-maps map-name forwarding-class class-name]` hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the `fragment-threshold` statement in the fragmentation map, the fragmentation threshold you set at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the `mrru` statement at the `[edit interfaces lsq-fpc/pic/port unit logical-unit-number]` hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from

1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 141.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. Because there is no MLPPP header, there is no sequence number information. Therefore, the software must take special measures to avoid packet reordering. To avoid packet reordering, the software places the packet on one of the *N* different T1 links. The link is determined by hashing the values in the header. For IP, the software computes the hash based on source address, destination address, and IP protocol. For MPLS, the software computes the hash based on up to five MPLS labels, or four MPLS labels and the IP header.

For UDP and TCP the software computes the hash based on the source and destination ports, as well as source and destination IP addresses. This guarantees that all packets belonging to the same TCP/UDP flow always pass through the same T1 link, and therefore cannot be reordered. However, it does not guarantee that the load on the various T1 links is balanced. If there are many flows, the load is usually balanced.

The *N* different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. If a packet has an MLPPP header, the sequence number field is used to put the packet back into sequence number order. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives and makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP

```
[edit interfaces]
lsq-1/1/0 {
    per-unit-scheduler;
    unit 0 {
        encapsulation multilink-ppp;
        mrru 2000;
        multilink-max-classes 4;
        family inet {
            address 20.1.1.1/24;
        }
        family mpls;
    }
}
ct1-1/1/4 {
    enable;
    no-partition interface-type t1;
```

```
}
t1-1/1/4 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
ct1-1/1/5 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/5 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
}
class-of-service {
    classifiers {
        inet-precedence myIPv4 {
            forwarding-class best-effort {
                loss-priority low code-points 000;
            }
            forwarding-class expedited-forwarding {
                loss-priority low code-points 001;
            }
            forwarding-class assured-forwarding {
                loss-priority low code-points 011;
            }
            forwarding-class network-control {
                loss-priority low code-points 111;
            }
        }
    }
    drop-profiles {
        dp1 {
            fill-level 50 drop-probability 0;
            fill-level 100 drop-probability 100;
```

```
        }
        dp2 {
            fill-level 50 drop-probability 0;
            fill-level 100 drop-probability 100;
        }
    }
    interfaces {
        lsq-1/1/0 {
            unit 0 {
                scheduler-map sm;
                fragmentation-map frag;
                rewrite-rules {
                    inet-precedence myRRIPv4;
                }
            }
        }
    }
    rewrite-rules {
        inet-precedence myRRIPv4 {
            forwarding-class best-effort {
                loss-priority low code-point 111;
            }
            forwarding-class expedited-forwarding {
                loss-priority low code-point 011;
            }
            forwarding-class assured-forwarding {
                loss-priority low code-point 001;
            }
            forwarding-class network-control {
                loss-priority low code-point 000;
            }
        }
    }
    scheduler-maps {
        sm {
            forwarding-class best-effort scheduler new;
            forwarding-class network-control scheduler new_nc;
            forwarding-class assured-forwarding scheduler new_af;
            forwarding-class expedited-forwarding scheduler new_ef;
        }
    }
    fragmentation-maps {
        frag {
```

```
            forwarding-class {
                best-effort {
                    multilink-class 3;
                }
                network-control {
                    multilink-class 0;
                }
                assured-forwarding {
                    multilink-class 2;
                }
                expedited-forwarding {
                    multilink-class 1;
                }
            }
        }
    }
    schedulers {
        new {
            transmit-rate 32k;
            shaping-rate 3m;
            priority low;
            drop-profile-map loss-priority low protocol any drop-profile dp1;
            drop-profile-map loss-priority high protocol any drop-profile dp2;
        }
        new_nc {
            transmit-rate 32k;
            shaping-rate 3m;
            priority strict-high;
        }
        new_af {
            transmit-rate 32k;
            shaping-rate 3m;
            priority medium-low;
        }
        new_ef {
            transmit-rate 32k;
            shaping-rate 3m;
            priority medium-high;
        }
    }
}
```

Understanding Multiclass MLPPP

Multiclass MLPPP makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, multiclass MLPPP allows different classes of traffic to have different latency guarantees. With multiclass MLPPP, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees. Multiclass MLPPP is defined in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. You can only configure multiclass MLPPP for link services intelligent queuing (LSQ) interfaces (`lsq-`) with MLPPP encapsulation.

Multiclass MLPPP greatly simplifies packet ordering issues that occur when multiple links are used. Without multiclass MLPPP, all voice traffic belonging to a single flow is hashed to a single link to avoid packet ordering issues. With multiclass MLPPP, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about voice services support on LSQ interfaces, see *Configuring Services Interfaces for Voice Services*.

If you do not configure multiclass MLPPP, fragments from different classes cannot be interleaved. All fragments for a single packet must be sent before the fragments from another packet are sent. Nonfragmented packets can be interleaved between fragments of another packet to reduce latency seen by nonfragmented packets. In effect, latency-sensitive traffic is encapsulated as regular PPP traffic, and bulk traffic is encapsulated as multilink traffic. This model works as long as there is a single class of latency-sensitive traffic, and there is no high-priority traffic that takes precedence over latency-sensitive traffic.

This approach to link fragmentation interleaving (LFI), used on the Link Services PIC, supports only two levels of traffic priority, which is not sufficient to carry the four to eight forwarding classes that are supported by M Series and T Series routers. For more information about the Link Services PIC support of LFI, see "Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces" on page 30.

NOTE: ACX Series routers do not support LFI.

Configuring both LFI and multiclass MLPPP on the same bundle is not necessary, nor is it supported, because multiclass MLPPP represents a superset of functionality. When you configure multiclass MLPPP, LFI is automatically enabled.

NOTE: The Junos OS implementation of multiclass MLPPP does not support compression of common header bytes, which is referred to in RFC 2686 as “prefix elision.”

Configuring Multiclass MLPPP on LSQ Interfaces

To configure multiclass MLPPP on a LSQ interface, you must specify how many multilink classes should be negotiated when a link joins the bundle, and you must specify the mapping of a forwarding class into an multiclass MLPPP class.

1. To specify how many multilink classes should be negotiated when a link joins the bundle, include the `multilink-max-classes` statement:

   ```
   multilink-max-classes number;
   ```

   You can include this statement at the following hierarchy levels:

   - `[edit interfaces interface-name unit logical-unit-number]`
   - `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]`

   The number of multilink classes can be 1 through 8. The number of multilink classes for each forwarding class must not exceed the number of multilink classes to be negotiated.

   NOTE: In ACX Series routers, the multilink classes can be 1 through 4.

2. To specify the mapping of a forwarding class into a multiclass MLPPP class, include the `multilink-class` statement at the `[edit class-of-service fragmentation-maps map-name forwarding-class class-name]` hierarchy level:

   ```
   [edit class-of-service fragmentation-maps map-name forwarding-class class-name]
   multilink-class number;
   ```

   The multilink class index number can be 0 through 7. The `multilink-class` statement and `no-fragmentation` statements are mutually exclusive.

   NOTE: In ACX Series routers, the multilink class index number can be 0 through 3. ACX Series routers do not support the `no-fragmentation` statement for fragmentation map.

3. To view the number of multilink classes negotiated, issue the `show interfaces lsq-fpc/pic/port.logical-unit-number detail` command.

Considerations for link services IQ (lsq) interfaces on ACX Series routers:

- The maximum number of multilink classes to be negotiated when a link joins the bundle that you can specify by using the `multilink-max-classes` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level is limited to 4.
- Fragmentation size is not specified under fragmentation map; instead, fragmentation size configured on the bundle is used.
- Compressed Real-Time Transport Protocol (RTP) is not supported.
- HDLC address and control field compression (ACFC) and PPP protocol field compression (PFC) are not supported.

SEE ALSO

Layer 2 Service Package Capabilities and Interfaces

Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP

Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP

Link Services Configuration for Junos Interfaces

RELATED DOCUMENTATION

Understanding Link and Multilink Services | 2

Configuring PPP

Examples: Bundling Multiple PPP Links on a Single Link Using MLPPP

IN THIS SECTION

- Example: Configuring a Multilink Interface with MLPPP | 55
- Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces | 56
- Example: Configuring an MLPPP Bundle | 58
- Example: Configuring a Link Services Interface with MLPPP | 64

- Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces | 65

Example: Configuring a Multilink Interface with MLPPP

```
[edit interfaces]
ml-1/0/0 {
    unit 1 {
        fragment-threshold 128;
        family inet {
            address 192.168.5.1/32 {
                destination 192.168.200.200;
            }
        }
    }
    unit 10 {
        family inet {
            address 10.1.1.3/32 {
                destination 10.1.1.2;
            }
        }
    }
}
t1-5/1/0 {
    unit 0 {
        family mlppp {
            bundle ml-1/0/0.1;
        }
    }
}
t1-5/1/1 {
    unit 0 {
        family mlppp {
            bundle ml-1/0/0.1;
        }
    }
}
t1-5/1/2 {
```

```
    unit 0 {
        family mlppp {
            bundle ml-1/0/0.1;
        }
    }
}
```

SEE ALSO

Link and Multilink Services Overview | 2

fragment-threshold

Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces

```
[edit interfaces]
at-0/0/0 {
    atm-options {
        pic-type atm2;
        vpi 10;
    }
    unit 0 {
        encapsulation atm-mlppp-llc;
        ppp-options {
            chap {
                access-profile pe-B-ppp-clients;
                local-name “pe-A-at-0/0/0”;
            }
        }
        keepalive interval 5 up-count 6 down-count 4;
        vci 10.120;
        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
}
at-0/0/1 {
    atm-options {
```

```
        pic-type atm2;
        vpi 11;
    }
    unit 1 {
        encapsulation atm-mlppp-llc;
        ppp-options {
            chap {
                access-profile pe-B-ppp-clients;
                local-name “ pe-A-at-0/0/0”;
            }
        }
        keepalive interval 5 up-count 6 down-count 4;
        vci 11.120;
        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
}
at-1/2/3 {
    atm-options {
        pic-type atm2;
        vpi 12;
    }
    unit 2 {
        encapsulation atm-mlppp-llc;
        ppp-options {
            chap {
                access-profile pe-B-ppp-clients;
                local-name “ pe-A-at-0/0/0”;
            }
        }
        keepalive interval 5 up-count 6 down-count 4;
        vci 12.120;
        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
}
...
ls-0/3/0 {
    encapsulation multilink-ppp;
    interleave-fragments;
    keepalive;
```

```
    unit 0 {
        mrru 4500;
        short-sequence;
        fragment-threshold 16320;
        drop-timeout 2000;
        encapsulation multilink-ppp;
        interleave-fragments;
        minimum-links 8;
        family inet {
            address 10.10.0.1/32 {
                destination 10.10.0.2;
            }
        }
        family iso;
        family inet6 {
            address 2001:DB8::0:1/32 {
                destination 2001:DB8::0:2;
            }
        }
    }
    ...
}
```

SEE ALSO

Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 136

Example: Configuring an MLPPP Bundle

IN THIS SECTION

- Requirements | 59
- Overview | 59
- Configuration | 59
- Verification | 63

This example shows how to configure an MLPPP bundle to increase traffic bandwidth.

Requirements

Before you begin, you should have two MX Series routers (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you create the MLPPP bundle lsq-0/0/0.0 at the logical unit level of the link services interface lsq-0/0/0 on the MX Series routers R0 and R1. You then add the two serial interfaces se-1/0/0 and se-1/0/1 as constituent links to the multilink bundle. In Figure 2 on page 59, your company's branch office is connected to its main branch using routers R0 and R1. You transmit data and voice traffic on two low-speed 1-Mbps serial links. To increase bandwidth, you configure MLPPP and join the two serial links se-1/0/0 and se-1/0/1 into the multilink bundle lsq-0/0/0.0. Then you configure LFI and CoS on R0 and R1 to enable them to transmit voice packets ahead of data packets.

Figure 2: Configuring MLPPP and LFI on Serial Links

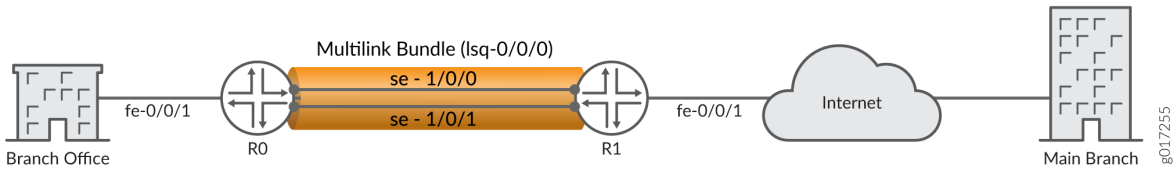

Configuration

IN THIS SECTION

- Procedure | **60**

Procedure

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]` hierarchy level.

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.10/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
set interfaces se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

```
For device R1
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.9/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure MLPPP bundle:

1. Create an interface on both the routers.

```
[edit]
user@host# edit interfaces lsq-0/0/0 unit 0
```

2. Configure a family inet and define the IP address on device R0.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.10/24
```

3. Configure a family inet and define the IP address on device R1.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.9/24
```

4. Specify the names of the constituent links to be added to the multilink bundle on both the routers.

```
[edit interfaces]
user@host# edit se-1/0/0 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0
[edit interfaces]
user@host# edit se-1/0/1 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0
```

5. Set the serial options to the same values for both interfaces on R0.

NOTE: R0 is set as a DCE device. The serial options are not set for interfaces on R1. You can set the serial options according to your network setup.

```
[edit interfaces]
user@host# set se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
user@host# set se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces lsq-0/0/0`, `show interfaces se-1/0/0`, and `show interfaces se-1/0/1` commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
    user@host# show interfaces lsq-0/0/0
family inet {
    address 10.0.0.10/24;
        }
    }
[edit]
```

```
user@host# show interfaces se-1/0/0
    clocking-mode dce;
        clock-rate 2.0mhz;
        }
        unit 0 {
            family mlppp {
            bundle lsq-0/0/0.0;
        }
    }
[edit]
    user@host# show interfaces se-1/0/1
serial-options {
    clocking-mode dce;
        clock-rate 2.0mhz;
        }
        unit 0 {
            family mlppp {
            bundle lsq-0/0/0.0;
        }
    }
```

```
For device R1
[edit]
user@host# show interfaces lsq-0/0/0
    family inet {
        address 10.0.0.9/24;
        }
    }
[edit]
    user@host# show interfaces se-1/0/0
    unit 0 {
    family mlppp {
        bundle lsq-0/0/0.0;
    }
}
[edit]
    user@host# show interfaces se-1/0/1
    unit 0 {
    family mlppp {
        bundle lsq-0/0/0.0;
```

```
    }
}
```

If you are done configuring the router, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying the MLPPP Bundle | **63**

Confirm that the configuration is working properly:

Verifying the MLPPP Bundle

Purpose

Verify that the constituent links are added to the bundle correctly.

Action

From operational mode, enter the `show interfaces lsq-0/0/0 statistics` command.

SEE ALSO

- Link and Multilink Services Overview | **2**
- Multilink Interfaces on Channelized MICs Overview | **7**
- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | **10**
- Example: Configuring Link Interfaces on Channelized MICs | **122**
- Example: Configuring Multilink Frame Relay FRF.15 | **96**
- Example: Configuring Multilink Frame Relay FRF.16 | **90**

Example: Configuring a Link Services Interface with MLPPP

```
[edit interfaces]
t1-0/0/0 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
}
t1-0/0/1 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
}
ls-0/3/0 {
    unit 0 {
        encapsulation multilink-ppp;
        family inet {
            address 10.16.1.2/32 {
                destination 10.16.1.1;
            }
        }
        family iso;
        family inet6 {
            address 2001:DB8::1:2/126;
        }
    }
}
```

SEE ALSO

Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 136

encapsulation (Logical Interface)

Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces

IN THIS SECTION

- Requirements | 65
- Overview | 65
- Configuration | 66
- Verification | 75

Inline Multilink PPP (*MLPPP*), Multlink Frame Relay (*FRF.16*), and Multilink Frame Relay End-to-End (*FRF.15*) for time-division multiplexing (*TDM*) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (*DPC*).

This example shows how to configure a Multilink PPP (*MLPPP*) bundle and Multilink Frame Relay End-to-End (*FRF.15*) for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1 (WAN interfaces).

Requirements

This example uses the following hardware and software components:

- Two MX Series Routers
- Junos OS Release 14.1 or later release

Before you begin, configure two MX Series routers (the MX240, MX480, or MX960) with at least two WAN interfaces that communicate over T1 links.

Overview

IN THIS SECTION

- Topology | 66

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (*LFI*) on the bundle, reducing high-priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

Configuring inline MLPPP for WAN interfaces benefits the following services:

- CE-PE link for Layer 3 VPN and DIA service with public switched telephone networks (PSTN)-based access networks
- PE-P link when PSTN is used for MPLS networks

In this example, to increase bandwidth, you configure MLPPP and join the T1 links into the multilink bundle. You aggregate T1 links to create the MLFR FRF.15 bundle on two MX Series routers, R0 and R1, and set the interface to `lsq-`. You configure logical units on the `lsq-` interface and set the family type to `inet` and an IP address. Then you configure an IP address for the multilink bundle on the unit level of the interface. You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be added to the multilink bundle and set the encapsulation type to `frame-relay`. You then define Router R0 as a DCE device and Router R1 as a DTE device. You set the DLCI value (range is from 16 through 1022). Finally, you set the multilink bundle to `lsq-`.

Topology

Figure 3: Configuring Inline MLPPP and Multilink Frame Relay End-End (FRF.15) for WAN Interfaces

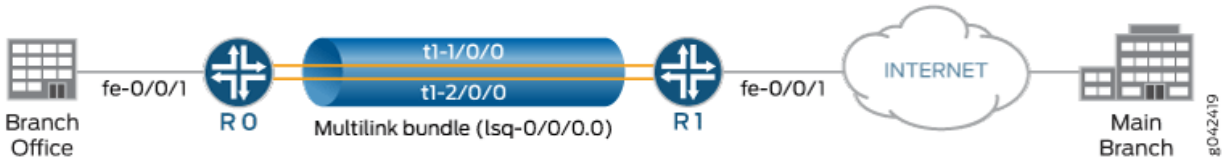

Configuration

IN THIS SECTION

CLI Quick Configuration | 67

- To Configure Router R0 | 68
- To Configure Router R1 | 69
- Results | 71

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]` hierarchy level.

Device R0

```
set chassis fpc 1 pic 0 multi-link-layer-2-inline
set interfaces lsq-1/0/0 unit 0 encapsulation multilink-ppp
set interfaces lsq-1/0/0 unit 0 family inet address 192.0.2.1/24
set interfaces lsq-1/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
set interfaces lsq-1/0/0 unit 1 family inet address 198.51.100.1/24
set interfaces t1-1/0/0:1 unit 0 family mlppp bundle lsq-1/0/0.0
set interfaces t1-1/0/0:2 unit 0 family mlppp bundle lsq-1/0/0.0
set interfaces t1-1/0/0:3 dce
set interfaces t1-1/0/0:4 dce
set interfaces t1-1/0/0:3 encapsulation frame-relay
set interfaces t1-1/0/0:4 encapsulation frame-relay
set interfaces t1-1/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-1/0/0.1
set interfaces t1-1/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-1/0/0.1
```

Device R1

```
set chassis fpc 2 pic 0 multi-link-layer-2-inline
set interfaces lsq-2/0/0 unit 0 encapsulation multilink-ppp
set interfaces lsq-2/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces lsq-2/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
set interfaces lsq-2/0/0 unit 1 family inet address 198.51.100.2/24
set interfaces t1-2/0/0:1 unit 0 family mlppp bundle lsq-2/0/0.0
set interfaces t1-2/0/0:2 unit 0 family mlppp bundle lsq-2/0/0.0
set interfaces t1-2/0/0:3 encapsulation frame-relay
set interfaces t1-2/0/0:4 encapsulation frame-relay
```

```
set interfaces t1-2/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-2/0/0.1
set interfaces t1-2/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-2/0/0.1
```

To Configure Router R0

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 CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces:

1. Enable inline Layer 2 bundling services.

```
[edit]
user@R0# set chassis fpc 1 pic 0 multi-link-layer-2-inline
```

2. Create the interface, specify a logical unit on the multilink bundle, and set the family type.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 0 family inet address 192.0.2.1/24
```

3. Specify the *encapsulation* type as MLPPP.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 0 encapsulation multilink-ppp
```

4. Create the interface, specify another logical unit on the multilink bundle, and set the family type.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 1 family inet address 198.51.100.1/24
```

5. Specify another unit and define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit]
user@R0# set interfaces lsq-1/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
```

6. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R0# set interfaces t1-1/0/0:1 unit 0 family mlppp bundle lsq-1/0/0.0
user@R0# set interfaces t1-1/0/0:2 unit 0 family mlppp bundle lsq-1/0/0.0
```

7. Define the router as a DCE device.

```
[edit]
user@R0# set interfaces t1-1/0/0:3 dce
user@R0# set interfaces t1-1/0/0:4 dce
```

8. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit ]
user@R0# set interfaces t1-1/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-1/0/0.1
user@R0# set interfaces t1-1/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-1/0/0.1
```

9. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R0# set interfaces t1-1/0/0:3 encapsulation frame-relay
user@R0# set interfaces t1-1/0/0:4 encapsulation frame-relay
```

To Configure Router R1

Step-by-Step Procedure

To configure inline MLPPP and Multilink Frame Relay End-to-End (FRF.15) for WAN Interfaces:

1. Enable inline Layer 2 bundling services.

```
[edit]
user@R1# set chassis fpc 2pic 0 multi-link-layer-2-inline
```

2. Create the interface, specify a logical unit on the multilink bundle and set the family type.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 0 family inet address 192.0.2.2/24
```

3. Specify the encapsulation type as MLPPP.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 0 encapsulation multilink-ppp
```

4. Create the interface, specify another logical unit on the multilink bundle and set the family type.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 1 family inet address 198.51.100.2/24
```

5. Specify another unit and define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit]
user@R1# set interfaces lsq-2/0/0 unit 1 encapsulation multilink-frame-relay-end-to-end
```

6. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R1# set interfaces t1-2/0/0:1 unit 0 family mlppp bundle lsq-2/0/0.0
user@R1# set interfaces t1-2/0/0:2 unit 0 family mlppp bundle lsq-2/0/0.0
```

7. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit ]
user@R1# set interfaces t1-2/0/0:3 unit 0 dlci 1 family mlfr-end-to-end bundle lsq-2/0/0.1
user@R1# set interfaces t1-2/0/0:4 unit 0 dlci 2 family mlfr-end-to-end bundle lsq-2/0/0.1
```

8. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit]
user@R1# set interfaces t1-2/0/0:3 encapsulation frame-relay
user@R1# set interfaces t1-2/0/0:4 encapsulation frame-relay
```

Results

For Router R0, from configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces lsq-1/0/0`, `show interfaces t1-1/0/0:1`, `show interfaces t1-1/0/0:2`, `show interfaces t1-1/0/0:3`, and `show interfaces t1-1/0/0:4` commands.

For Router R1, from configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces lsq-2/0/0`, `show interfaces t1-2/0/0:1`, `show interfaces t1-2/0/0:2`, `show interfaces t1-2/0/0:3`, and `show interfaces t1-2/0/0:4` commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

For Router R0:

```
[edit]
user@R0# show chassis
fpc 1 {
    pic 0 {
        multi-link-layer-2-inline;
    }
}
```

```
[edit]
user@R0# show interfaces lsq-1/0/0
unit 0 {
    encapsulation multilink-ppp;
    family inet {
        address 192.0.2.1/24;
    }
}
unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
```

```
        address 198.51.100.1/24;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:1
unit 0 {
    family mlppp {
        bundle lsq-1/0/0.0;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:2
unit 0 {
    family mlppp {
        bundle lsq-1/0/0.0;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:3
dce;
encapsulation frame-relay;
unit 0 {
    dlci 1;
    family mlfr-end-to-end {
        bundle lsq-1/0/0.1;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:4
dce;
encapsulation frame-relay;
unit 0 {
    dlci 2;
```

```
    family mlfr-end-to-end {
        bundle lsq-1/0/0.1;
    }
}
```

If you are done configuring the router, enter `commit` from configuration mode.

For Router R1:

```
[edit]
user@R1# show chassis
fpc 2{
    pic 0 {
        multi-link-layer-2-inline;
    }
}
```

```
[edit]
user@R1# show interfaces lsq-2/0/0
unit 0 {
    encapsulation multilink-ppp;
    family inet {
        address 192.0.2.2/24;
    }
}
unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
        address 198.51.100.2/24;
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:1
unit 0 {
    family mlppp {
        bundle lsq-2/0/0.0;
```

```
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:2
unit 0 {
    family mlppp {
        bundle lsq-2/0/0.0;
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:3
encapsulation frame-relay;
unit 0 {
    dlci 1;
    family mlfr-end-to-end {
        bundle lsq-2/0/0.1;
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:4
encapsulation frame-relay;
unit 0 {
    dlci 2;
    family mlfr-end-to-end {
        bundle lsq-2/0/0.1;
    }
}
```

If you are done configuring the router, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying the MLPPP Bundle and the MLFR FRF.15 Configuration | 75

Verifying the MLPPP Bundle and the MLFR FRF.15 Configuration

Purpose

Verify that the constituent links are added to the bundle correctly.

Action

From operational mode, run the `show interfaces lsq-1/0/0 extensive` command.

Sample Output

command-name

```
user@R0> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0, Enabled, Physical link is Up
  Interface index: 292, SNMP ifIndex: 1065, Generation: 4986
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
```

```
  Frame exceptions:
    Oversized frames                  0
    Errored input frames              0
    Input on disabled link/bundle     0
    Output for disabled link/bundle   0
    Queuing drops                     0
  Buffering exceptions:
    Packet data buffer overflow       0
    Fragment data buffer overflow     0
  Assembly exceptions:
    Fragment timeout                  0
    Missing sequence number           0
    Out-of-order sequence number      0
    Out-of-range sequence number      0
  Hardware errors (sticky):
    Data memory error                 0
    Control memory error              0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control

  Logical interface lsq-1/0/0.0 (Index 327) (SNMP ifIndex 113518) (Generation 6213)
    Flags: Hardware-Down Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Last flapped: 2014-04-24 04:37:39 PDT (00:08:50 ago)
    Bandwidth: 0
    Bundle links information:
      Active bundle links       0
      Removed bundle links      2
      Disabled bundle links     0
    Bundle options:
      MRRU                              1504
      Remote MRRU                       N/A
      Drop timer period                 32767
      Inner PPP Protocol field compression enabled
      Sequence number format            long (24 bits)
```

```
    Fragmentation threshold              0
    Links needed to sustain bundle       1
    Multilink classes                    0
    Link layer overhead                  4.0 %
  Multilink class 0 status:
    Received sequence number        0x0
    Transmit sequence number        0xffffffff
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
    Multilink class drop timeout    0 (ms)
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  IPV6 Transit Statistics         Packets        Bytes
    Network:
      Input :                      0              0
      Output:                      0              0
  Link:
    t1-1/0/0:1.0
      Up time: 00:00:00
      Input :             0          0             0            0
      Output:             0          0             0            0
    t1-1/0/0:2.0
      Up time: 00:00:00
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink detail statistics:
  Bundle:
    Fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
```

```
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      LFI:
        Input :             0          0             0            0
        Output:             0          0             0            0
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured, mpls: Not-
configured
    Protocol inet, MTU: 1500, Generation: 6263, Route table: 0
      Flags: Sendbcast-pkt-to-re, Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: Unspecified, Generation: 4211

  Logical interface lsq-1/0/0.1 (Index 328) (SNMP ifIndex 113519) (Generation 6214)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
    Last flapped: 2014-04-24 04:46:00 PDT (00:00:29 ago)
    Bandwidth: 3072kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Bundle options:
      MRRU                                1504
      Drop timer period                   32767
      Inner PPP Protocol field compression enabled
      Sequence number format              short (12 bits)
      Fragmentation threshold             0
      Links needed to sustain bundle      1
      Multilink classes                   0
      Link layer overhead                 4.0 %
    Multilink class 0 status:
      Received sequence number        0x0
      Transmit sequence number        0xffffffff
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
      Multilink class drop timeout    0 (ms)
```

```
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      t1-1/0/0:3.0
        Up time: 00:00:29
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-1/0/0:4.0
        Up time: 00:00:29
        Input :             0          0             0            0
        Output:             0          0             0            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      LFI:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500, Generation: 6264, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100/24, Local: 198.51.100.1, Broadcast: Unspecified, Generation: 4213
```

From the operational mode, enter the `show interfaces lsq-2/0/0 extensive` command.

```
user@R1> show interfaces lsq-2/0/0 extensive
Physical interface: lsq-2/0/0, Enabled, Physical link is Up
  Interface index: 262, SNMP ifIndex: 44421, Generation: 270
  Encapsulation: Multilink-PPPLink-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Last flapped   : Never
```

```
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Frame exceptions:
    Oversized frames                 0
    Errored input frames             0
    Input on disabled link/bundle    0
    Output for disabled link/bundle  0
    Queuing drops                    0
  Buffering exceptions:
    Packet data buffer overflow      0
    Fragment data buffer overflow    0
  Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
    Out-of-range sequence number     0
  Hardware errors (sticky):
    Data memory error                0
    Control memory error             0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                   0                   0
    1                                0                   0                   0
    2                                0                   0                   0
    3                                0                   0                   0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control

  Logical interface lsq-2/0/0.0 (Index 354) (SNMP ifIndex 44422) (Generation 167)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPPEncapsulation:
Multilink-PPP
```

```
    Last flapped: 2014-04-24 04:50:19 PDT (00:00:51 ago)
    Bandwidth: 3072kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Bundle options:
      MRRU                                1504
      Remote MRRU                         1504
      Drop timer period                   32767
      Inner PPP Protocol field compression enabled
      Sequence number format              long (24 bits)
      Fragmentation threshold             0
      Links needed to sustain bundle      1
      Multilink classes                   0
      Link layer overhead                 4.0 %
    Multilink class 0 status:
      Received sequence number        0x0
      Transmit sequence number        0xffffffff
     Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
      Multilink class drop timeout    0 (ms)
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    IPV6 Transit Statistics         Packets        Bytes
      Network:
        Input :                      0              0
        Output:                      0              0
    Link:
      t1-2/0/0:1.0
```

```
        Up time: 00:00:51
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-2/0/0:2.0
        Up time: 00:00:48
        Input :             0          0             0            0
        Output:             0          0             0            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      LFI:
        Input :             0          0             0            0
        Output:             0          0             0            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
    Protocol inet, MTU: 1500, Generation: 199, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.2, Broadcast: Unspecified, Generation: 153

  Logical interface lsq-4/0/0.1 (Index 355) (SNMP ifIndex 44423) (Generation 168)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
    Last flapped: 2014-04-24 04:50:19 PDT (00:00:51 ago)
    Bandwidth: 3072kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Bundle options:
      MRRU                                1504
      Drop timer period                   32767
      Inner PPP Protocol field compression enabled
      Sequence number format              short (12 bits)
      Fragmentation threshold             0
      Links needed to sustain bundle      1
      Multilink classes                   0
      Link layer overhead                 4.0 %
    Multilink class 0 status:
      Received sequence number        0x0
```

```
    Transmit sequence number        0xffffffff
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
    Multilink class drop timeout    0 (ms)
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Link:
    t1-2/0/0:3.0
      Up time: 00:00:51
      Input :             0          0             0            0
      Output:             0          0             0            0
    t1-2/0/0:4.0
      Up time: 00:00:51
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink detail statistics:
  Bundle:
    Fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Non-fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
    LFI:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet, MTU: 1500, Generation: 200, Route table: 0
    Flags: Sendbcast-pkt-to-re
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100/24, Local: 198.51.100.2, Broadcast: Unspecified, Generation: 155
```

From operational mode, enter the `show interfaces lsq-1/0/0 statistics` command.

```
user@R0> show interfaces lsq-1/0/0 statistics
Physical interface: lsq-1/0/0, Enabled, Physical link is Up
  Interface index: 292, SNMP ifIndex: 1065
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface lsq-1/0/0.0 (Index 327) (SNMP ifIndex 113518)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Last flapped: 2014-04-24 04:50:19 PDT (00:01:59 ago)
    Bandwidth: 3072kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    IPV6 Transit Statistics         Packets        Bytes
      Network:
        Input :                      0              0
        Output:                      0              0
    Link:
      t1-1/0/0:1.0
        Up time: 00:01:59
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-1/0/0:2.0
```

```
      Up time: 00:01:56
      Input :             0          0             0            0
      Output:             0          0             0            0
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.0.2/24, Local: 192.0.2.1

Logical interface lsq-1/0/0.1 (Index 328) (SNMP ifIndex 113519)
  Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
  Last flapped: 2014-04-24 04:50:29 PDT (00:01:49 ago)
  Bandwidth: 3072kbps
  Bundle links information:
    Active bundle links       2
    Removed bundle links      0
    Disabled bundle links     0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Link:
    t1-1/0/0:3.0
      Up time: 00:01:49
      Input :             0          0             0            0
      Output:             0          0             0            0
    t1-1/0/0:4.0
      Up time: 00:01:49
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 198.51.100/24, Local: 198.51.100.1
```

From operational mode, enter the `show interfaces lsq-2/0/0 statistics` command.

```
user@R1> show interfaces lsq-2/0/0 statistics
Physical interface: lsq-2/0/0, Enabled, Physical link is Up
  Interface index: 262, SNMP ifIndex: 44421
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface lsq-2/0/0.0 (Index 354) (SNMP ifIndex 44422)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Last flapped: 2014-04-24 04:50:19 PDT (00:04:33 ago)
    Bandwidth: 3072kbps
    Bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    IPV6 Transit Statistics         Packets        Bytes
      Network:
        Input :                      0              0
        Output:                      0              0
    Link:
      t1-2/0/0:1.0
        Up time: 00:04:33
        Input :             0          0             0            0
        Output:             0          0             0            0
     t1-2/0/0:2.0
        Up time: 00:04:30
        Input :             0          0             0            0
        Output:             0          0             0            0
```

```
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.0.2/24, Local: 192.0.2.2

Logical interface lsq-2/0/0.1 (Index 355) (SNMP ifIndex 44423)
  Flags: Up Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
  Last flapped: 2014-04-24 04:50:19 PDT (00:04:33 ago)
  Bandwidth: 3072kbps
  Bundle links information:
    Active bundle links       2
    Removed bundle links      0
    Disabled bundle links     0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Link:
    t1-2/0/0:3.0
      Up time: 00:04:33
      Input :             0          0             0            0
      Output:             0          0             0            0
    t1-2/0/0:4.0
      Up time: 00:04:33
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 198.51.100/24, Local: 198.51.100.2
```

SEE ALSO

Inline MLPPP for WAN Interfaces Overview

Enabling MLPPP Link Fragmentation and Interleaving | 24

Example: Configuring Multilink Frame Relay FRF.15 | 96

mlfr-uni-nni-bundles-inline

multi-link-layer-2-inline

pic (MX Series Routers)

show interfaces (Link Services IQ)

Examples: Bundling Multiple Frame Relay DLCIs into a Single Link Using MLFR

IN THIS SECTION

- Example: Configuring a Multilink Interface with MLFR FRF.15 | 88
- Example: Configuring Multilink Frame Relay FRF.16 | 90
- Example: Configuring Multilink Frame Relay FRF.15 | 96
- Configuring DLCIs on Link Services Logical Interfaces | 100
- Example: Configuring a Link Services PIC with MLFR FRF.16 | 102
- Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces | 103

Example: Configuring a Multilink Interface with MLFR FRF.15

```
[edit interfaces]
ml-1/0/0 {
    unit 1 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 192.168.5.2/32 {
                destination 192.168.5.3;
            }
        }
    }
    unit 10 {
```

```
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.1.1.3/32 {
                destination 10.1.1.2;
            }
        }
    }
}
t1-5/1/0 {
    unit 0 {
        dlci 16;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.1;
        }
    }
}
t1-5/1/1 {
    unit 0 {
        dlci 17;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.10;
        }
    }
}
t1-5/1/2 {
    unit 0 {
        dlci 26;
        encapsulation multilink-frame-relay-end-to-end;
        family mlfr-end-to-end {
            bundle ml-1/0/0.10;
        }
    }
}
```

SEE ALSO

Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 136

encapsulation (Logical Interface)

Example: Configuring Multilink Frame Relay FRF.16

IN THIS SECTION

- Requirements | 90
- Overview | 90
- Configuration | 90
- Verification | 95

This example shows how to configure MLFR FRF.16 for additional bandwidth, load balancing, and redundancy.

Requirements

Before you begin, you should have two MX Series 5G Universal Routing Platforms configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 interfaces to create an MLFR FRF.16 bundle on two MX Series, R0 and R1. You configure the chassis interface and specify the number of MLFR FRF.16 bundles to be created on the interface. You then specify the channel to be configured as a multilink bundle and create interface lsq-0/0/0:0. You set the multilink bundle as an MLFR FRF.16 bundle by specifying the MLFR UNI NNI encapsulation type.

Then you define R0 as a DCE device and R1 as a DTE device. You configure a logical unit on the multilink bundle lsq-0/0/0:0, and set the family type to inet. You then assign a DLCI of 400 and an IP address of 10.0.0.10/24 to the multilink bundle. You create the T1 interfaces, t1-2/0/0 and t1-2/0/1, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to lsq-0/0/0:0.

Configuration

IN THIS SECTION

- Procedure | 91

Procedure

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]` hierarchy level.

```
For device R0
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.10/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
For device R1
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.9/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure an MLFR FRF.16 bundle:

1. Configure a chassis interface.

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

2. Specify the number of MLFR bundles.

```
[edit chassis]
user@host# set fpc 0 pic 0 mlfr-uni-nni-bundles 1
```

3. Create an interface.

```
[edit]
user@host# edit interfaces lsq-0/0/0:0
```

4. Specify the MLFR encapsulation type.

```
[edit interfaces lsq-0/0/0:0]
user@host# set encapsulation multilink-frame-relay-uni-nni
```

5. Set the router R0 as a DCE device.

```
[edit]
user@host# edit interfaces lsq-0/0/0
user@host# set dce
```

6. Specify a logical unit on the multilink bundle and set the family type.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 400 family inet address 10.0.0.10/24
```

7. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
user@host# set t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
```

8. Specify the multilink bundle to which the interface is to be added as a constituent link on device R0.

```
[edit interfaces t1-2/0/0]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

9. Specify the multilink bundle to which the interface is to be added as a constituent link on device R1.

```
[edit interfaces t1-2/0/1]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Results

From configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces lsq-0/0/0`, `show interfaces lsq-0/0/0:0`, `show interfaces t1-2/0/0`, and `show interfaces t1-2/0/1` commands for the routers R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
    user@host# show chassis
    fpc 0 {
    pic 0 {
    mlfr-uni-nni-bundles 1;
        }
    }
    [edit]
        user@host# show interfaces lsq-0/0/0
        dce;
    unit 0 {
    dlci 400;
        family inet {
        address 10.0.0.10/24;
    }
    }
    [edit]
        user@host# show interfaces lsq-0/0/0:0
    encapsulation multilink-frame-relay-uni-nni;
    [edit]
        user@host# show interfaces t1-2/0/0
```

```
        encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
    family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
    }
}
[edit]
    user@host# show interfaces t1-2/0/1
    encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
    bundle lsq-0/0/0:0;
}
}
```

```
For device R1
[edit]
    user@host# show chassis
    unit 0 {
    dlci 400;
    family inet {
        address 10.0.0.9/24;
    }
    }
    [edit]
        user@host# show interfaces lsq-0/0/0:0
    encapsulation multilink-frame-relay-uni-nni;
    [edit]
        user@host# show interfaces t1-2/0/0
        encapsulation multilink-frame-relay-uni-nni;
        unit 0 {
            family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
    }
    }
    [edit]
        user@host# show interfaces t1-2/0/1
        encapsulation multilink-frame-relay-uni-nni;
        unit 0 {
            family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
```

```
    }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying the MLFR FRF.16 Configuration | 95

Confirm that the configuration is working properly:

Verifying the MLFR FRF.16 Configuration

Purpose

Verify the MLFR FRF.16 configuration.

Action

From operational mode, enter the `show interfaces` command.

SEE ALSO

Link and Multilink Services Overview | 2

Multilink Interfaces on Channelized MICs Overview | 7

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 10

Example: Configuring Link Interfaces on Channelized MICs | 122

Example: Configuring an MLPPP Bundle | 58

Example: Configuring Multilink Frame Relay FRF.15

IN THIS SECTION

- Requirements | 96
- Overview | 96
- Configuration | 96
- Verification | 100

This example shows how to configure MLFR FRF.15 for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1, E1, and serial links.

Requirements

Before you begin, you should have two MX Series 5G Universal Routing Platforms (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 links to create the MLFR FRF.15 bundle on two MX Series routers, R0 and R1, and set the interface to lsq-0/0/0. You configure a logical unit on the lsq-0/0/0 interface and set the family type to `inet` with address 10.0.0.4/24. Then you configure an IP address for the multilink bundle on the unit level of the interface.

You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be added to the multilink bundle as t1-2/0/0 and t1-2/0/1 and set the encapsulation type to `frame-relay`. You then define R0 as a DCE device and R1 as a DTE device. You set the DLCI value to 100 (range is from 16 through 1022). Finally, you set the multilink bundle to lsq-0/0/0.0.

Configuration

IN THIS SECTION

- Procedure | 97

Procedure

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]` hierarchy level.

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.4/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

```
For device R1
 set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.5/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure the MLFR FRF.15 bundle:

1. Create an interface on both the routers.

```
[edit]
user@host# edit interfaces lsq-0/0/0 unit 0
```

2. Set a logical unit on the interface and define the family type for the routers R0 and R1.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.4/24
user@host# set family inet address 10.0.0.5/24
```

3. Define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set encapsulation multilink-frame-relay-end-to-end
```

4. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation frame-relay
user@host# set t1-2/0/1 encapsulation frame-relay
```

5. Define the router R0 as a DCE device.

```
[edit interfaces]
user@host# edit lsq-0/0/0
user@host# set dce
```

6. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces lsq-0/0/0`, `show interfaces t1-2/0/0`, and `show interfaces t1-2/0/1` commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
```

```
    user@host# show interfaces lsq-0/0/0
    dce;
unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
        dlci 100;
    family inet {
        address 10.0.0.4/24;
            }
        family mlfr-end-to-end {
        bundle lsq-0/0/0.0;
            }
        }
[edit]
    user@host#show interfaces t1-2/0/0
encapsulation frame-relay;
    [edit]
        user@host# show interfaces t1-2/0/1
    encapsulation frame-relay;
```

```
For device R1
[edit]
    user@host# show interfaces lsq-0/0/0
    unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
        dlci 100;
        family inet {
    address 10.0.0.5/24;
    }
    family mlfr-end-to-end {
        bundle lsq-0/0/0.0;
    }
}
[edit]
    user@host# show interfaces t1-2/0/0
encapsulation frame-relay;
    [edit]
        user@host# show interfaces t1-2/0/1
    encapsulation frame-relay;
```

If you are done configuring the router, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying the MLFR FRF.15 Configuration | 100

Confirm that the configuration is working properly:

Verifying the MLFR FRF.15 Configuration

Purpose

Verify the MLFR FRF.15 configuration.

Action

From operational mode, enter the `show interfaces` command.

SEE ALSO

Link and Multilink Services Overview | 2

Multilink Interfaces on Channelized MICs Overview | 7

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 10

Example: Configuring Link Interfaces on Channelized MICs | 122

Example: Configuring an MLPPP Bundle | 58

Configuring DLCIs on Link Services Logical Interfaces

IN THIS SECTION

- Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles | 101
- Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles | 101

For link services interfaces only, you can configure multiple DLCIs for MLFR FRF.16 or MLPPP bundles.

DLCIs are not supported on multilink interfaces.

Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles

For link services interfaces only, you can configure multiple point-to-point DLCIs for each MLFR FRF.16 or MLPPP bundle. A channelized interface, such as `ls-1/1/1:0`, denotes a single MLFR FRF.16 bundle. To configure a DLCI, include the `dlci` statement:

```
dlci dlci-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The DLCI identifier is a value from 16 through 1022. Numbers 1 through 15 are reserved for future use.

When you configure point-to-point connections, the maximum transmission unit (MTU) sizes on both sides of the connection must be the same.

Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles

For link services interfaces only, you can configure multiple multicast-capable DLCIs for each MLFR FRF.16 bundle. A channelized interface, such as `ls-1/1/1:0`, denotes a single MLFR FRF.16 bundle. By default, Frame Relay connections assume unicast traffic. If your Frame Relay switch performs multicast replication, you can configure the link services connection to support multicast traffic by including the `multicast-dlci` statement:

```
multicast-dlci dlci-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The DLCI identifier is a value from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

You can configure multicast support only on point-to-multipoint link services connections. Multicast-capable DLCIs are not supported on multilink interfaces.

If keepalives are enabled, causing the interface to send Local Management Interface (LMI) messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. For more information, see "Configuring Keepalives on Link Services Physical Interfaces" on page 134.

SEE ALSO

Link and Multilink Services Overview | 2

Configuring Link and Multilink Services Logical Interfaces | 134

Example: Configuring a Link Services PIC with MLFR FRF.16

```
[edit chassis]
fpc 1 {
    pic 2 {
        mlfr-uni-nni-bundles 5;
    }
}
[edit interfaces]
t1-0/0/0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/2/0:0;
        }
    }
}
t1-0/0/1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/2/0:0;
        }
    }
}
ls-1/2/0:0 {
    dce;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
```

```
        dlci 26;
        family inet {
            address 10.26.1.1/32 {
                destination 10.26.1.2;
            }
        }
    }
}
```

SEE ALSO

Configuring Link Services Physical Interfaces | **134**

encapsulation (Physical Interface)

Example: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces

IN THIS SECTION

- Requirements | **104**
- Overview | **104**
- Configuration | **105**
- Verification | **111**

Inline Multilink PPP (*MLPPP*), Multilink Frame Relay (*FRF.16*), and Multilink Frame Relay End-to-End (*FRF.15*) for time-division multiplexing (*TDM*) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (*DPC*).

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (*LFI*) on the bundle, reducing high priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

This example shows how to configure Multilink Frame Relay (*FRF.16*) for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1 (WAN interfaces).

Requirements

This example uses the following hardware and software components:

- Two MX Series Routers
- Junos OS Release 14.1 or later release

Before you begin, configure two MX Series routers (the MX240, MX480, or MX960) with at least two WAN interfaces that communicate over T1 links.

Overview

IN THIS SECTION

- Topology | **104**

In this example, you aggregate T1 interfaces to create an MFR FRF.16 bundle on two MX Series routers, R0 and R1. You configure the chassis interface and specify the number of MFR FRF.16 bundles to be created on the interface. You then specify the channel to be configured as a multilink bundle and create interface `lsq-`. You set the multilink bundle as an MFR FRF.16 bundle by specifying the `multilink-frame-relay-uni-nni` encapsulation type. Then you define Router R0 as a DCE device and Router R1 as a DTE device. You configure a logical unit on the multilink bundle `lsq-`, and set the family type to `inet`. You create the T1 interfaces, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to `lsq-`.

Topology

Figure 4: Configuring Inline Multilink Frame Relay (FRF.16) for WAN Interfaces

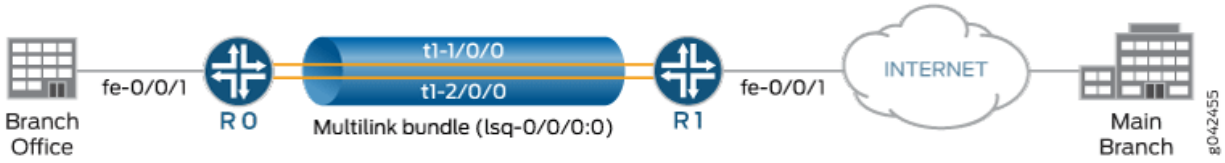

Configuration

IN THIS SECTION

- CLI Quick Configuration | 105
- To Configure Router R0 | 106
- To Configure Router R1 | 107
- Results | 108

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]` hierarchy level.

Device R0

```
set chassis fpc 1 pic 0 mlfr-uni-nni-bundles-inline 1
set interfaces lsq-1/0/0:0 dce
set interfaces lsq-1/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-1/0/0:0 unit 0 dlci 10
set interfaces lsq-1/0/0:0 unit 1 dlci 20
set interfaces lsq-1/0/0:0 unit 0 family inet address 10.1.1.1/24
set interfaces lsq-1/0/0:0 unit 1 family inet address 192.0.2.1./24
set interfaces t1-1/0/0:5 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-1/0/0:6 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-1/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
set interfaces t1-1/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
```

Device R1

```
set chassis fpc 2 pic 0 mlfr-uni-nni-bundles-inline 1
set interfaces lsq-2/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-2/0/0:0 unit 0 dlci 10
set interfaces lsq-2/0/0:0 unit 1 dlci 20
set interfaces lsq-2/0/0:0 unit 0 family inet address 10.1.1.2/24
set interfaces lsq-2/0/0:0 unit 1 family inet address 192.0.2.2/24
set interfaces t1-2/0/0:5 encapsulation multilink-frame-relay-uni-nni
```

```
set interfaces t1-2/0/0:6 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
set interfaces t1-2/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
```

To Configure Router R0

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 CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure inline Multlink Frame Relay (*FRF.16*) for WAN Interfaces:

1. Configure a chassis interface and specify the number of MFR bundles.

```
[edit]
user@R0# set chassis fpc 1 pic 0 mlfr-uni-nni-bundles-inline 1
```

2. Create the interface and specify the MFR encapsulation type.

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 encapsulation multilink-frame-relay-uni-nni
```

3. Set Router R0 as a DCE device.

```
[edit ]
user@R0# set interfaces lsq-1/0/0:0 dce
```

4. Specify the DLCI value.

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 unit 0 dlci 10
user@R0# set interfaces lsq-1/0/0:0 unit 1 dlci 20
```

5. Specify a logical unit on the multilink bundle and set the family type.

```
[edit]
user@R0# set interfaces lsq-1/0/0:0 unit 0 family inet address 10.1.1.1/24
user@R0# set interfaces lsq-1/0/0:0 unit 1 family inet address 192.0.2.1/24
```

6. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit]
user@R0# set interfaces t1-1/0/0:5 encapsulation multilink-frame-relay-uni-nni
user@R0# set interfaces t1-1/0/0:6 encapsulation multilink-frame-relay-uni-nni
```

7. Specify the multilink bundle to which the interface is to be added as a constituent link on Router R0.

```
[edit]
user@R0# set interfaces t1-1/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
user@R0# set interfaces t1-1/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-1/0/0:0
```

To Configure Router R1

Step-by-Step Procedure

To configure inline Multilink Frame Relay (*FRF.16*) for WAN Interfaces:

1. Configure a chassis interface and specify the number of MFR bundles.

```
[edit]
user@R1# set chassis fpc 2 pic 0 mlfr-uni-nni-bundles-inline 1
```

2. Create the interface and specify the MFR encapsulation type.

```
[edit]
user@R1# set interfaces lsq-2/0/0:0 encapsulation multilink-frame-relay-uni-nni
```

3. Specify the DLCI value.

```
[edit]
user@R0# set interfaces lsq-2/0/0:0 unit 0 dlci 10
user@R0# set interfaces lsq-2/0/0:0 unit 1 dlci 20
```

4. Specify a logical unit on the multilink bundle and set the family type.

```
[edit]
user@R0# set interfaces lsq-2/0/0:0 unit 0 family inet address 10.1.1.2/24
user@R0# set interfaces lsq-2/0/0:0 unit 1 family inet address 192.0.2.2/24
```

5. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit]
user@R1# set interfaces t1-2/0/0:5 encapsulation multilink-frame-relay-uni-nni
user@R1# set interfaces t1-2/0/0:6 encapsulation multilink-frame-relay-uni-nni
```

6. Specify the multilink bundle to which the interface is to be added as a constituent link on Router R1.

```
[edit]
user@R1# set interfaces t1-2/0/0:5 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
user@R1# set interfaces t1-2/0/0:6 unit 0 family mlfr-uni-nni bundle lsq-2/0/0:0
```

Results

For Router R0, from configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces lsq-1/0/0:0`, `show interfaces t1-1/0/0:5`, and `show interfaces t1-1/0/0:6` commands.

For Router R1, from configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces lsq-2/0/0:0`, `show interfaces t1-2/0/0:5`, and `show interfaces t1-2/0/0:6` commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

For Router R0:

```
[edit]
user@R0# show chassis
```

```
fpc 1 {
    pic 0 {
        mlfr-uni-nni-bundles-inline 1;
    }
}
```

```
[edit]
user@R0# show interfaces lsq-1/0/0:0
dce;
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    dlci 10;
    family inet {
        address 10.1.1.1/24;
    }
}
unit 1 {
    dlci 20;
    family inet {
        address 192.0.2.1/24;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:5
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-1/0/0:0;
    }
}
```

```
[edit]
user@R0# show interfaces t1-1/0/0:6
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-1/0/0:0;
```

```
    }
}
```

If you are done configuring the router, enter `commit` from configuration mode.

For Router R1:

```
[edit]
user@R1# show chassis
fpc 2{
    pic 0 {
        mlfr-uni-nni-bundles-inline 1;
    }
}
```

```
[edit]
user@R1# show interfaces lsq-2/0/0:0
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    dlci 10;
    family inet {
        address 10.1.1.2/24;
    }
}
unit 1 {
    dlci 20;
    family inet {
        address 192.0.2.2/24;
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:5
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-2/0/0:0;
```

```
    }
}
```

```
[edit]
user@R1# show interfaces t1-2/0/0:6
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-2/0/0:0;
    }
}
```

If you are done configuring the router, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- Verifying the MFR FRF.16 Configuration | **111**

Verifying the MFR FRF.16 Configuration

Purpose

Verify the MFR FRF.16 configuration.

Action

From operational mode, run the `show interfaces lsq-1/0/0:0 extensive` command.

Sample Output

command-name

```
user@R0> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
  Interface index: 261, SNMP ifIndex: 122042, Generation: 4955
```

```
Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : Never
Statistics last cleared: Never
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type                       DCE
  MRRU                              1508
  Bandwidth                         3072kbps
  Fragmentation threshold           0
  Red differential delay limit      120
  Yellow differential delay limit   72
  Red differential delay action     Remove link
 Reassembly drop timer             65535
  Links needed to sustain bundle    1
  Link layer overhead               4.0 %
  LIP Hello timer                   10
      Acknowledgement timer         4
      Acknowledgement retries       2
  Bundle class                      A
  LMI type                          Q.933 Annex A
    T391 LIV polling timer          10
    T392 polling verification timer 15
    N391 full status polling count  6
    N392 error threshold            3
    N393 monitored event count      4
Q.933 Annex A LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
LMI statistics:
  Input : 52 (last seen 00:00:01 ago)
  Output: 54 (last sent 00:00:01 ago)
DTE statistics:
  Enquiries sent                    : 0
  Full enquiries sent               : 0
  Enquiry responses received        : 0
  Full enquiry responses received   : 0
DCE statistics:
  Enquiries received                : 44
  Full enquiries received           : 8
  Enquiry responses sent            : 46
  Full enquiry responses sent       : 8
Common statistics:
  Unknown messages received         : 0
```

```
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  Interface transmit statistics: Disabled
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Multilink Frame Relay UNI NNI bundle errors:
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Exception events                0
  Multilink Frame Relay UNI NNI bundle statistics:
                    Frames        fps         Bytes          bps

    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink Frame Relay UNI NNI bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
  Multilink Frame Relay UNI NNI active bundle links statistics:
                    Frames        fps         Bytes          bps
    t1-1/0/0:5
      Up time: 00:08:18
      Input :             0          0             0            0
      Output:             0          0             0            0
      Current differential delay       0.1 ms
      Recent high differential delay   0.8 ms
      Times over red diff delay        0
      Times over yellow diff delay     0
      LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
```

```
      Rcv:      2       1       0       50      49       0       0
      Xmt:     16       2       0       49      50       1       0
    t1-1/0/0:6
      Up time: 00:08:18
      Input :              0          0             0            0
      Output:              0          0             0            0
      Current differential delay       0.0 ms
      Recent high differential delay   0.7 ms
      Times over red diff delay        0
      Times over yellow diff delay     0
      LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
      Rcv:      2       1       0       50      49       0       0
      Xmt:     16       2       0       49      50       1       0

  Logical interface lsq-1/0/0:0.0 (Index 336) (SNMP ifIndex 122044) (Generation 6209)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
    Last flapped: 2014-04-24 04:13:05 PDT (00:08:18 ago)
    Multilink class 0 status:
      Received sequence number        0x0
      Transmit sequence number        0xffffffff
Packet drops                      0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
      Multilink class drop timeout    0 (ms)
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      t1-1/0/0:5
        Up time: 00:08:18
        Input :             0          0             0            0
        Output:             0          0             0            0
```

```
      t1-1/0/0:6
        Up time: 00:08:18
        Input :             0          0             0            0
        Output:             0          0             0            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500, Generation: 6258, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: Unspecified, Generation: 4209
    DLCI 10
      Flags: Active
      Total down time: 01:15:17 sec, Last down: 01:23:28 ago
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0

  Logical interface lsq-1/0/0:0.1 (Index 337) (SNMP ifIndex 122067) (Generation 6210)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
    Last flapped: 2014-04-24 04:13:05 PDT (00:08:18 ago)
    Multilink class 0 status:
      Received sequence number        0x0
      Transmit sequence number        0xffffffff
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
      Multilink class drop timeout    0 (ms)
    Statistics         Frames        fps         Bytes          bps
    Bundle:
```

```
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      t1-1/0/0:5
        Up time: 00:08:18
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-1/0/0:6
        Up time: 00:08:18
        Input :             0          0             0            0
        Output:             0          0             0            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500, Generation: 6260, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: Unspecified, Generation: 4207
    DLCI 20
      Flags: Active
      Total down time: 01:15:17 sec, Last down: 01:23:28 ago
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
  DLCI statistics:
    Active DLCI  :2  Inactive DLCI  :0
```

From the operational mode, enter the `show interfaces lsq-2/0/0:0 extensive` command.

```
user@R1> show interfaces lsq-2/0/0:0 extensive
Physical interface: lsq-2/0/0:0, Enabled, Physical link is Up
```

```
Interface index: 232, SNMP ifIndex: 44389, Generation: 235
Link-level type: Multilink-FR-UNI-NNI,MTU: 1508
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Last flapped   : Never
Statistics last cleared: Never
Hold-times     : Up 0 ms, Down 0 ms
Multilink Frame Relay UNI NNI bundle options:
  Device type                      DTE
  MRRU                             1508
  Bandwidth                        3072kbps
  Fragmentation threshold          0
  Red differential delay limit     120
  Yellow differential delay limit  72
  Red differential delay action    Remove link
  Reassembly drop timer            65535
  Links needed to sustain bundle   1
  Link layer overhead              4.0 %
  LIP Hello timer                  10
      Acknowledgement timer        4
      Acknowledgement retries      2
  Bundle class                     A
  LMI type                         Q.933 Annex A
    T391 LIV polling timer         10
    T392 polling verification timer 15
    N391 full status polling count  6
    N392 error threshold           3
    N393 monitored event count     4
Q.933 Annex A LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
LMI statistics:
  Input : 80 (last seen 00:00:10 ago)
  Output: 100 (last sent 00:00:10 ago)
DTE statistics:
  Enquiries sent                     : 82
  Full enquiries sent                : 16
  Enquiry responses received         : 67
  Full enquiry responses received    : 13
DCE statistics:
  Enquiries received                 : 0
  Full enquiries received            : 0
  Enquiry responses sent             : 0
  Full enquiry responses sent        : 0
Common statistics:
```

```
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  Interface transmit statistics: Disabled
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Multilink Frame Relay UNI NNI bundle errors:
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Exception events                0
  Multilink Frame Relay UNI NNI bundle statistics:
                    Frames        fps         Bytes          bps

    Multilink:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Network:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink Frame Relay UNI NNI bundle links information:
      Active bundle links       2
      Removed bundle links      0
      Disabled bundle links     0
  Multilink Frame Relay UNI NNI active bundle links statistics:
                    Frames        fps         Bytes          bps
    t1-2/0/0:5
      Up time: 00:12:57
      Input :             0          0             0            0
      Output:             0          0             0            0
      Current differential delay       0.0 ms
      Recent high differential delay   2.8 ms
      Times over red diff delay        0
      Times over yellow diff delay     0
```

```
      LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
      Rcv:      1       2       0       77      78       0       0
      Xmt:     14       1       0       78      77       0       0
    t1-2/0/0:6
      Up time: 00:12:57
      Input :             0          0             0            0
      Output:             0          0             0            0
      Current differential delay       0.0 ms
      Recent high differential delay   2.8 ms
      Times over red diff delay        0
      Times over yellow diff delay     0
      LIP:add_lnk lnk_ack lnk_rej    hello hel_ack lnk_rem rem_ack
      Rcv:      1       2       0       77      78       0       0
      Xmt:     14       1       0       78      77       0       0

  Logical interface lsq-2/0/0:0.0 (Index 348) (SNMP ifIndex 44399) (Generation 161)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
    Last flapped: 2014-04-24 04:13:05 PDT (00:12:57 ago)
    Multilink class 0 status:
      Received sequence number        0x0
      Transmit sequence number        0xffffffff
      Packet drops                    0 (0 bytes)
      Fragment drops                  0 (0 bytes)
      MRRU exceeded                   0
      Fragment timeout                0
      Missing sequence number         0
      Out-of-order sequence number    0
      Out-of-range sequence number    0
      Packet data buffer overflow     0
      Fragment data buffer overflow   0
      Multilink class drop timeout    0 (ms)
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      t1-2/0/0:5
        Up time: 00:12:57
        Input :             0          0             0            0
```

```
      Output:             0          0             0            0
   t1-2/0/0:6
      Up time: 00:12:57
      Input :             0          0             0            0
      Output:             0          0             0            0
  Multilink detail statistics:
  Bundle:
    Fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
    Non-fragments:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet, MTU: 1500, Generation: 193, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.2, Broadcast: Unspecified, Generation: 149
  DLCI 10
    Flags: Active, DCE-Configured
    Total down time: 00:03:18 sec, Last down: 00:15:38 ago
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0

Logical interface lsq-2/0/0:0.1 (Index 349) (SNMP ifIndex 44400) (Generation 162)
  Flags: Up Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
  Last flapped: 2014-04-24 04:13:05 PDT (00:12:57 ago)
  Multilink class 0 status:
    Received sequence number        0x0
    Transmit sequence number        0xffffffff
    Packet drops                    0 (0 bytes)
    Fragment drops                  0 (0 bytes)
    MRRU exceeded                   0
    Fragment timeout                0
    Missing sequence number         0
    Out-of-order sequence number    0
    Out-of-range sequence number    0
    Packet data buffer overflow     0
    Fragment data buffer overflow   0
    Multilink class drop timeout    0 (ms)
  Statistics         Frames        fps         Bytes          bps
```

```
    Bundle:
      Multilink:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Network:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      t1-2/0/0:5
        Up time: 00:12:57
        Input :             0          0             0            0
        Output:             0          0             0            0
     t1-2/0/0:6
        Up time: 00:12:57
        Input :             0          0             0            0
        Output:             0          0             0            0
    Multilink detail statistics:
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
      Non-fragments:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500, Generation: 194, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.2, Broadcast: Unspecified, Generation: 151
    DLCI 20
      Flags: Active, DCE-Configured
      Total down time: 00:03:18 sec, Last down: 00:15:38 ago
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
  DLCI statistics:
    Active DLCI  :2  Inactive DLCI  :0
```

SEE ALSO

Inline MLPPP for WAN Interfaces Overview

Enabling MLPPP Link Fragmentation and Interleaving | 24

mlfr-uni-nni-bundles-inline

multi-link-layer-2-inline

pic (MX Series Routers)

show interfaces (Link Services IQ)

Example: Configuring Link Interfaces on Channelized MICs

IN THIS SECTION

- Requirements | 122
- Overview | 122
- Configuration on SDH OC3/STM1 (Multi-Rate) MIC with SFP | 123
- Verification | 128

Requirements

This example uses the following hardware and software components:

- MX Series routers that support Link Interfaces on Channelized MICs with Junos OS

Overview

This example provides information about configuring the link interfaces on the following channelized MICs:

- SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)

- /SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)

You need to first partition each port on the MICs to configure the link interfaces and then you configure the link interfaces for bundles. An MLPPP bundle involves "bundling" multiple interfaces into a single, logical interface that uses only one IP address. For more information about MLPPP bundles, see "Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links" on page 10. .

For more information about multilink-based protocols on MX series routers with Multiservices DPC, see "Multilink Interfaces on Channelized MICs Overview" on page 7.

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

NOTE: You can set the values for each parameter according to your requirement. The values given in this example are for illustration purposes only.

Configuration on SDH OC3/STM1 (Multi-Rate) MIC with SFP

IN THIS SECTION

- CLI Quick Configuration | **123**
- Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces | **125**

To partition each port on the MIC and configure the link interfaces on it and to configure the link interfaces for bundles, perform the following tasks:

CLI Quick Configuration

To quickly configure synchronization on the aforementioned routers, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

[edit]

```
set interfaces coc12-5/2/0 partition 1 interface-type coc1
set interfaces coc12-5/2/0 partition 1 oc-slice 1
set interfaces coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
```

```
set interfaces coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
set interfaces coc1-5/2/0:1 no-partition interface-type t3
set interfaces coc1-5/2/0:3 no-partition interface-type t3
set interfaces coc1-5/2/0:2 partition 1 interface-type ct1
set interfaces coc1-5/2/0:2 partition 2 interface-type t1
set interfaces coc1-5/2/0:2 partition 3 interface-type ct1
set interfaces coc1-5/2/0:2 partition 4 interface-type t1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
set interfaces ct3-2/0/0 no-partition interface-type t3
set interfaces ct3-2/0/0 partition 1 interface-type t1
set interfaces ct3-2/0/0 partition 2 interface-type t1
set interfaces ct3-2/0/0 partition 3 interface-type t1
set interfaces ct3-2/0/0 partition 4 interface-type ct1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces t1-5/2/0:2:2 unit 0 family mlppp bundle rlsq0.1
set interfaces ds-5/2/0:2:1:1 unit 0 family mlppp bundle rlsq0.0
set interfaces ds-5/2/0:2:1:2 unit 0 family mlppp bundle rlsq0.0
set interfaces t3-5/2/0:3 unit 0 family mlppp bundle rlsq0.2
set interfaces t3-5/2/0:1 unit 0 family mlppp bundle rlsq0.2
set interfaces t1-5/2/0:2:2 encapsulation multilink-frame-relay-uni-nni unit 0 family mlfr-uni-nni bundle
rlsq0:0
set interfaces t1-5/2/0:2:4 encapsulation multilink-frame-relay-uni-nni unit 0 family mlfr-uni-nni bundle
rlsq0:0
set interfaces ds-5/2/0:2:1:1 encapsulation multilink-frame-relay-uni-nni unit 0 family mlfr-uni-nni
bundle rlsq0:0
set interfaces ds-5/2/0:2:1:2 encapsulation multilink-frame-relay-uni-nni unit 0 family mlfr-uni-nni
bundle rlsq0:0
set interfaces t1-5/2/0:2:2 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end bundle
rlsq0.0
set interfaces t1-5/2/0:2:4 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end bundle
rlsq0.0
set interfaces ds-5/2/0:2:1:1 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end bundle
rlsq0.0
set interfaces ds-5/2/0:2:1:2 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end bundle
rlsq0.0
set interfaces t3-5/2/0:1 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end bundle rlsq0.1
set interfaces t3-5/2/0:3 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end bundle rlsq0.1
```

Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces

Step-by-Step Procedure

To partition each port on the channelized MICs:

1. Configure the coc12-5/2/0 interface by setting the partition option to 1 with the sublevel interface type set to coc1.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 1 interface-type coc1
```

2. Configure the coc12-5/2/0 interface with the OC-slice range (OC-slice range specifies the bandwidth size required for the interface type you are configuring) set to 1.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 1 oc-slice 1
```

3. Configure the coc12-5/2/0 interface by setting the partition option to 2 with the sublevel interface type set to coc1 and the OC-slice range set to 2.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
```

4. Configure the coc12-5/2/0 interface by setting the partition option to 3 with the sublevel interface type set to coc1 and the OC-slice range set to 3.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
```

5. Configure the coc1-5/2/0:1 interface as a clear channel by setting the no-partition option for the sublevel interface type t3 . (A clear channel consolidates the entire bandwidth of a channelized interface into a single unpartioned stream that looks like a standard interface.)

```
[edit interfaces]
user@host# set coc1-5/2/0:1 no-partition interface-type t3
```

6. Configure the coc1-5/2/0:3 interface as a clear channel by setting the no-partition option for the sublevel interface type t3.

```
[edit interfaces]
user@host# set coc1-5/2/0:3 no-partition interface-type t3
```

7. Configure the coc1-5/2/0:2 interface by setting the partition option to 1 and 3 with the sublevel interface type set to ct1. Configure the coc1-5/2/0:2 interface by setting the partition option to 2 and 4 with the sublevel interface type set to t1.

```
[edit interfaces]
user@host# set coc1-5/2/0:2 partition 1 interface-type ct1
user@host# set coc1-5/2/0:2 partition 2 interface-type t1
user@host# set coc1-5/2/0:2 partition 3 interface-type ct1
user@host# set coc1-5/2/0:2 partition 4 interface-type t1
```

8. Configure the ct1-5/2/0:2:1 interface by setting the partition option to 1 and 2 with the sublevel interface type set to ds. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```

9. Configure a clear channel on the channelized interface ct3-2/0/0 by setting the no-partition option to the sublevel interface type t3 (a clear channel consolidates the entire bandwidth of a channelized interface into a single unpartitioned stream that looks like a standard interface).

```
[edit interfaces]
user@host# set ct3-2/0/0 no-partition interface-type t3
```

10. Configure a clear channel on the channelized interface ct3-2/0/0 by setting the partition option to 1, 2, and 3 with the sublevel interface type ds. Configure the ct3-2/0/0 interface by setting the partition option to 4 with the sublevel interface type ct1.

```
[edit interfaces]
user@host# set ct3-2/0/0 partition 1 interface-type t1
user@host# set ct3-2/0/0 partition 2 interface-type t1
```

```
user@host# set ct3-2/0/0 partition 3 interface-type t1
user@host# set ct3-2/0/0 partition 4 interface-type ct1
```

11. Configure the ct1-2/0/0:4 interface by setting the partition option to 1 and 2 with the sublevel interface type set to ds. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```

Results

Display the results of partitioning each port on the MIC and configuring the link interfaces :

Results for CHOC12/3 interfaces

```
user@host# show interfaces
coc12-5/2/0 {
    partition 1 oc-slice 1 interface-type coc1;
    partition 2 oc-slice 2 interface-type coc1;
    partition 3 oc-slice 3 interface-type coc1;
}
coc1-5/2/0:1 {
    no-partition interface-type t3;
}
coc1-5/2/0:3 {
    no-partition interface-type t3;
}
coc1-5/2/0:2 {
    partition 1 interface-type ct1;
    partition 2 interface-type t1;
    partition 3 interface-type ct1;
    partition 4 interface-type t1;
}
```

Verification

IN THIS SECTION

- Verifying the MLPPP Bundle | 128
- Verifying the MLFR FRF.15 Configuration | 128
- Verifying the MLFR FRF.16 Configuration | 129

Confirm that the configuration is working properly.

Verifying the MLPPP Bundle

Purpose

Verify that the constituent links are added to the bundle correctly.

Action

From operational mode, enter the `show interfaces lsq-`*`fpc/pic/port`* command.

Meaning

The output displays the constituent links that are added to the bundle. For more information about the `show interfaces lsq-`*`fpc/pic/port`* operational command, see the CLI Explorer.

Verifying the MLFR FRF.15 Configuration

Purpose

Verify the MLFR FRF.15 configuration.

Action

From operational mode, enter the `show interfaces lsq-`*`fpc/pic/port`* command.

Meaning

The output displays the standard status information about the specified link services IQ interface. For more information about the `show interfaces lsq-`*`fpc/pic/port`* operational command, see the CLI Explorer.

Verifying the MLFR FRF.16 Configuration

Purpose

Verify the MLFR FRF.16 configuration.

Action

From operational mode, enter the `show interfaces lsq-`*`fpc/pic/port`* command.

Meaning

The output displays the standard status information about the specified link services IQ interface. For more information about the `show interfaces lsq-`*`fpc/pic/port`* operational command, see the CLI Explorer.

RELATED DOCUMENTATION

- Link and Multilink Services Overview | 2
- Multilink Interfaces on Channelized MICs Overview | 7
- Example: Configuring an MLPPP Bundle | 58
- Example: Configuring Multilink Frame Relay FRF.15 | 96
- Example: Configuring Multilink Frame Relay FRF.16 | 90

Unified ISSU on Inline LSQ Interfaces Overview

On MX Series routers that support Unified in-service software upgrade (ISSU), the feature is supported on inline link services intelligent queuing (LSQ) interfaces . Unified ISSU enables an upgrade between two Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Inline Multilink PPP (MLPPP) for WAN interfaces provides bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC). The inline LSQ logical interface (lsq-*slot/pic/*0) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide bundling services for Layer 2 packets that do not need a services PIC.

Unified ISSU support for inline LSQ interfaces provides backward compatibility with Junos OS releases in which this support is not available. You can perform a unified ISSU process between a Junos OS release that supports unified ISSU for inline LSQ interfaces and a Junos OS release in which unified ISSU support for inline LSQ interfaces is not available. This backward compatibility does not apply to the scenario where part of the software (such as the kernel software) is upgraded while another part (such as the Packet Forwarding Engine software) is not upgraded. Unified ISSU infrastructure provides APIs to store persistent data, perform cold boot (without resetting power), and manage post-ISSU reboot connectivity and synchronization with Routing Engine kernel and other processes. Unified ISSU is also supported for redundant LSQ (rlsq-) interfaces configured in hot-standby and warm-standby mode.

The following line cards on MX Series routers that support unified ISSU for inline LSQ interfaces:

- 8-port Channelized SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 4-port Channelized SDH OC3/STM1 (Multi-Rate) MIC with SFP1 (MIC-3D-4CHOC3-2CHOC12)
- MPC Type 1 3D Q (MX-MPC1-3D-Q)
- MPC Type 2 3D Q (MX-MPC2-3D-Q)

LSQ interfaces support the following three types of multilink data traffic:

- Complete multilink packets (for which bandwidth and encapsulation are set)—These packets do not require the reassembly memory.
- Fragmented multilink packets–These packets require reassembly memory to form a complete packets until all the fragments of all packets arrive.
- Link fragment interleaving (LFI) packets—These packets do not have any multilink header.

There are two time intervals or durations called dark windows during the unified ISSU process. These dark windows denote the brief traffic disruption and outage in handling of packets that can occur when the Packet Forwarding Engines restart owing to the upgraded software being installed on them. The duration of these traffic disconnections vary in length based on several factors, such as the platform and the Junos OS version being upgraded. The following two dark windows are observed:

- The first dark window is during the unified ISSU boot phase. During this time period, host-bound packets are affected because the new microkernel image is being initialized and the path to the host is not discovered yet. The aging calculations for reassembly and fragmentation are stopped, which causes the fragmented traffic types to be dropped. However, LFI packets and complete multilink packets are not disrupted.
- The second dark window occurs during the hardware synchronization phase. During this time period, the hardware learns all the new states from the software. ASICs reinitialize themselves with the new configurations. During this period, all kinds of traffic drops are permissible. The duration for this dark window is approximately 4–5 seconds. During the second dark window, because packets are

dropped at the fragmentation side, a receive-ml-seq number mismatch occurs at the other end of the connection (reassembly or receiver side). This behavior can cause additional packet drops.

During a unified ISSU, a switchover from the primary LSQ interface to the secondary lsq interface does not occur and the bundle remains up. The impact on fragmented packets (to be reassembled) remains the same as the packets on LSQ interfaces. LFI and non-fragmented packets are traversed until the second dark window starts. The dark window duration for rlsq- interfaces is the same as that for LSQ interfaces. No particular attribute is maintained only on the local MPC hardware. All of the QoS functionalities behave in the same manner as the pre-unified ISSU phase.

During the unified ISSU process, none of the multilink counters are updated (except the multilink sequence number). The counters are reset and incremented from zero after unified ISSU is completed. The receive-ml-seq and send ml-seq counters need to be updated for reassembly to take place properly but the update of these counters is not exported to the Packet Forwarding Engine microkernel or the Routing Engine.

It is assumed that bundle links do not flap during the unified ISSU window. The following operations occur at the fragmentation and reassembly ends during a unified ISSU process:

When unified ISSU is completed, the send multilink sequence (send-ml-seq) number starts from zero and it causes the receive multilink sequence (receive-ml-seq) number mismatch at the other end of the connection, resulting in additional packet drops. Consider the case where LSQ bundles and links can be hosted on different Packet Forwarding Engines in a dual Packet Forwarding Engine setup. In this scenario, it is possible that these two Packet Forwarding Engines are in different states of the ISSU process, where dark windows for both of them might follow different timelines. This condition adds to more drops.

At the reassembly side:

1. The reassembly is stopped at the unified ISSU preparation or initialization phase, which can cause a longer dark window for fragmented packets. At the time of unified ISSU, reassembly memory is used for unified ISSU processing and is not available, starting from the unified ISSU preparation phase up to the unified ISSU hardware synchronization stage.
2. The two activities of link-load calculations and aging of fragments might impact the traversal of packets during unified ISSU.

At the fragmentation side:

1. .The ml-seq number counter, which is needed for multilink fragmentation to work, is not part of the unified ISSU counters.
2. The host-generated packets (that are Layer 2-rewrite packets) are transferred as part of the fabric header and saved as binary large objects (BLOBs).
3. For LFI traffic, hash calculation must work properly for the links to be load-balanced. If the media link state changes during unified ISSU, the bundle state might also be affected.

RELATED DOCUMENTATION

Enabling Inline LSQ Services

CHAPTER 3

Configuring the Physical and Logical Interfaces in a Link and Multilink Configuration

IN THIS CHAPTER

- Configuring Link Services Physical Interfaces | 134
- Configuring Link and Multilink Services Logical Interfaces | 134

Configuring Link Services Physical Interfaces

You configure link services interface properties at the logical unit and physical interface level. Default settings for link services physical interface properties are described in "Configuring Link Services Physical Interfaces" on page 134.

The following sections explain how to configure link services physical interfaces:

For information about link services physical interface properties that can also be configured at the logical unit level, see "Configuring Link and Multilink Services Logical Interfaces" on page 134.

RELATED DOCUMENTATION

Link and Multilink Services Overview | 2

Example: Configuring a Link Services Interface with Two Links | 14

Configuring Link and Multilink Services Logical Interfaces

IN THIS SECTION

- Multilink and Link Services Logical Interface Configuration Overview | 135
- Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 136
- Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces | 137
- Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces | 139
- Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces | 140
- Configuring MRRU on Multilink and Link Services Logical Interfaces | 141
- Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces | 142
- Configuring CoS on Link Services Interfaces | 143

Multilink and Link Services Logical Interface Configuration Overview

IN THIS SECTION

- Default Settings for Multilink and Link Services Logical Interfaces | 135

You configure multilink and link services interface properties at the logical unit level. Default settings for multilink and link services logical interface properties are described in "Default Settings for Multilink and Link Services Logical Interfaces" on page 135.

For general information about logical unit properties or `family inet` properties, see the Junos OS Network Interfaces Library for Routing Devices. For information about multilink and link services properties you configure at the `family inet` hierarchy level, see "Configuring the Links in a Multilink or Link Services Bundle" on page 12.

NOTE: The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about logical interface properties, see the Junos OS Network Interfaces Library for Routing Devices.

Default Settings for Multilink and Link Services Logical Interfaces

Table 4 on page 135 lists the default settings for multilink and link services statements, together with the other permitted values or value ranges.

Table 4: Multilink and Link Services Logical Interface Statements

Option	Default Value	Possible Values
DLCI	None	16 through 1022
Drop timeout period	500 ms for bundles greater than or equal to the T1 bandwidth value and 1500 ms for other bundles.	0 through 2000 milliseconds

Table 4: Multilink and Link Services Logical Interface Statements *(Continued)*

Option	Default Value	Possible Values
Encapsulation	For multilink interfaces, `multilink-ppp`. For link services interfaces, `multilink-frame-relay-end-to-end`.	`multilink-frame-relay-end-to-end`, `multilink-ppp`
Fragmentation threshold	0 bytes	128 through 16,320 bytes (*N*x64)
Interleave fragments	disabled	enabled, disabled
Minimum links	1 link	1 through 8 links
Maximum received reconstructed unit (MRRU)	1504 bytes	1500 through 4500 bytes
Sequence ID format for MLPPP	24 bits	12 or 24 bits
Sequence ID format for MLFR FRF.15 and FRF.16	12 bits	12 bits

See "Default Settings for Link Services Interfaces" on page 134 for statements that apply to link services physical interfaces only.

SEE ALSO

Configuring DLCIs on Link Services Logical Interfaces | 100

Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces | 30

Configuring Encapsulation for Multilink and Link Services Logical Interfaces

Multilink and link services interfaces support MLPP logical interface encapsulation type.

By default, the logical interface encapsulation type on multilink interfaces is MLPPP. The default logical interface encapsulation type on link services interfaces is MLFR end-to-end. For general information on encapsulation, see the Junos OS Network Interfaces Library for Routing Devices.

You can also configure physical interface encapsulation on link services interfaces. For more information, see "Configuring Encapsulation for Link Services Physical Interfaces" on page 134.

To configure multilink or link services encapsulation, include the `encapsulation` *type* statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

SEE ALSO

> *encapsulation (Logical Interface)*

Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces

By default, the drop timeout parameter is disabled. You can configure a drop timeout value to provide a recovery mechanism if individual links in the multilink or link services bundle drop one or more packets. Drop timeout is not a differential delay tolerance setting, and does not limit the overall latency. However, you need to make sure the value you set is larger than the expected differential delay across the links, so that the timeout period does not elapse under normal jitter conditions, but only when there is actual packet loss. You can configure differential delay tolerance for link services interfaces only. For more information, see "Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16" on page 134.

To configure the drop timeout value, include the `drop-timeout` statement:

```
drop-timeout milliseconds;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For link services interfaces, you also can configure the drop timeout value at the physical interface level by including the `drop-timeout` statement at the `[edit interfaces ls-`*`fpc`*`/`*`pic`*`/`*`port`*`:`*`channel`* `mlfr-uni-nni-bundle-options]` hierarchy level:

```
drop-timeout milliseconds;
```

By default, the drop timer has a value of 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Values can range from 1 through 2000 milliseconds. Values less than 5 milliseconds are not recommended, and a configured value of 0 reverts to the default value of 2000 milliseconds.

NOTE: For multilink or link services interfaces, if a packet or fragment encounters an error condition and is destined for a disabled bundle or link, it does not contribute to the dropped packet and frame counts in the per-bundle statistics. The packet is counted under the global error statistics and is not included in the global output bytes and output packet counts. This unusual accounting happens only if the error conditions are generated inside the multilink interface, not if the packet encounters errors on the wire or elsewhere in the network.

If you configure the `drop-timeout` statement with a value of `0`, it disables any resequencing by the PIC for the specified class of MLPPP traffic. Packets are forwarded with the assumption that they arrived in sequence, and forwarding of fragmented packets is disabled for all classes. Fragments dropped as a result of this setting will increment the counter at the class level.

Alternatively, you can configure the `drop-timeout` statement at the `[edit class-of-service fragmentation-maps` *`map-name`* `forwarding-class` *`class`*`]` hierarchy level. The behavior and the default and range values are identical, but the setting applies only to the specified forwarding class. Configuration at the bundle level overrides configuration at the class-of-service level.

By default, compression of the inner PPP header in the MLPPP payload is enabled. To disable compression, include the `disable-mlppp-inner-ppp-pfc` statement at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]` hierarchy level. For example:

```
interfaces lsq-1/2/0 {
    unit 0 {
        encapsulation multilink-ppp;
        disable-mlppp-inner-ppp-pfc;
        multilink-max-classes 4;
        family inet {
            address 10.50.1.2/30;
        }
```

```
    }
}
```

For more information about CoS configuration, see the Junos OS Class of Service User Guide for Routing Devices. You can view the configured drop-timeout value and the status of inner PPP header compression by issuing the `show interfaces` *`interface-name`* `extensive` command.

Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces

For multilink and link services logical interfaces with MLPPP encapsulation only, you can configure a *fragmentation threshold* to limit the size of packet payloads transmitted across the individual links within the multilink circuit. The software splits any incoming packet that exceeds the fragmentation threshold into smaller units suitable for the circuit size; it reassembles the fragments at the other end, but does not affect the output traffic stream. The threshold value affects the payload only; it does not affect the MLPPP header. By default, the fragmentation threshold parameter is disabled.

NOTE: To ensure proper load balancing:

- For Link Services MLFR (FRF.15 and FRF.16) interfaces, do not include the `fragment-threshold` statement in the configuration.
- For MLPPP interfaces, do not include both the `fragment-threshold` statement and the `short-sequence` statement in the configuration.
- For MLFR (FRF.15 and FRF.16) and MLPPP interfaces, if the MTU of links in a bundle is less than the bundle MTU plus encapsulation overhead, then fragmentation is automatically enabled. You should avoid this situation for MLFR (FRF.15 and FRF.16) interfaces and for MLPPP interfaces on which short-sequencing is enabled.

To configure a fragmentation threshold value, include the `fragment-threshold` statement:

```
fragment-threshold bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For link services interfaces, you also can configure a fragmentation threshold value at the physical interface level by including the `fragment-threshold` statement at the `[edit interfaces ls-`*`fpc`*`/`*`pic`*`/`*`port`*`:`*`channel`* `mlfr-uni-nni-bundle-options]` hierarchy level:

```
fragment-threshold bytes;
```

The maximum fragment size can be from 128 through 16,320 bytes. The Junos OS automatically subdivides packet payloads that exceed this value. Any value you set must be a multiple of 64 bytes (*N*x64). The default value, 0, results in no fragmentation.

SEE ALSO

> *fragment-threshold*

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

You can set the minimum number of links that must be up for the multilink bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The `minimum-links` value should be identical on both ends of the bundle.

To set the minimum number, include the `minimum-links` statement:

```
minimum-links number;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]`
- `[edit logical-systems` *`logical-system-name`* `interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]`

For link services interfaces, you also can configure the minimum number of links at the physical interface level by including the `minimum-links` statement at the `[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]` hierarchy level:

```
minimum-links number;
```

The number can be from 1 through 8. The maximum number of links supported in a bundle is 8. When 8 is specified, all configured links of a bundle must be up.

Configuring MRRU on Multilink and Link Services Logical Interfaces

The *maximum received reconstructed unit* (*MRRU*) is the maximum packet size that the multilink interface can process. It is similar to a maximum transmission unit (MTU), but applies only to multilink bundles. By default, the MRRU is set to 1500 bytes. You can configure a different MRRU value if the peer equipment allows this. The MRRU accounts for the original payload, for example the Layer 3 protocol payload, but does not include the 2-byte PPP header or the additional MLPPP or MLFR header applied while the individual multilink packets are traversing separate links in the bundle.

To configure a different MRRU value, include the `mrru` statement:

```
mrru bytes;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]`

NOTE: ACX Series routers do not support logical systems.

For link services interfaces, you also can configure a different MRRU at the physical interface level by including the `mrru bytes` statement at the `[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]` hierarchy level. The MRRU size can range from 1500 through 4500 bytes.

NOTE: If you set the MRRU on a bundle to a value larger than the MTU of the individual links within it, you must enable a fragmentation threshold for that bundle. Set the threshold to a value no larger than the smallest MTU of any link included in the bundle.

Determine the appropriate MTU size for the bundle by ensuring that the MTU size does not exceed the sum of the encapsulation overhead and the MTU sizes for the links in the bundle.

You can configure separate `family mtu` values on the following protocol families under bundle interfaces: `inet`, `inet6`, `iso`, and `mpls`. If not configured, the default value of 1500 is used on all except for `mpls` configurations, in which the value 1488 is used.

NOTE: ACX Series routers do not support `family inet6` on MLPPP interfaces.

NOTE: The effective family MTU might be different from the MTU value specified for MLPPP configurations, because it is adjusted downward by the remote MRRU's constraints.

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces

For MLPPP, the sequence header format is set to 24 bits by default. You can configure an alternative value of 12 bits, but 24 bits is considered the more robust value for most networks.

To configure a different sequence header value, include the `short-sequence` statement:

```
short-sequence;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]`
- `[edit logical-systems` *`logical-system-name`* `interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]`

For MLFR FRF.15, the sequence header format is set to 24 bits by default. This is the only valid option.

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

Configuring CoS on Link Services Interfaces

IN THIS SECTION

- Example: Configuring CoS on Link Services Interfaces | 143

For link services IQ (lsq-) interfaces, Junos OS class of service (CoS) is fully supported and functions as described in the Junos OS Class of Service User Guide for Routing Devices. For more information and detailed configuration examples, see *Layer 2 Service Package Capabilities and Interfaces*.

On SRX Series Firewalls, the lsq- interface is an internal interface, which is not associated with a physical interface. For information about link services on SRX Series Firewalls, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

Example: Configuring CoS on Link Services Interfaces

Configure CoS on a link services interface and its constituent link interfaces.

Packets that do not match the firewall filters are sent to a queue that performs load balancing by sending fragments to all constituent links.

Packets that match the firewall filters are sent to a queue that does not support packet fragmentation and reassembly; instead, this traffic is load-balanced by sending each packet flow to a different constituent link. Each packet that matches a firewall filter is subjected to a hash on the IP source address and the IP destination address to determine the packet flow to which each packet belongs.

When you configure the MLPPP encapsulation type routing protocol packets smaller than 128 bytes are sent to the network-control queue on the constituent link interface. This keeps routing protocols operating normally, even when low-speed links are congested by regular packets.

```
[edit interfaces]
ls-7/0/0 {
    unit 0 {
        encapsulation multilink-ppp;
        interleave-fragments;
        family inet {
            filter {
                output lfi_ls_filter;
            }
```

```
            address 10.54.0.2/32 {
                destination 10.54.0.1;
            }
        }
    }
}
ge-7/2/0 {
    unit 0 {
        family inet {
            address 192.168.1.1/24;
        }
    }
}
ce1-7/3/6 {
    no-partition interface-type e1;
}
e1-7/3/6 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle ls-7/0/0.0;
        }
    }
}
ce1-7/3/7 {
    no-partition interface-type e1;
}
e1-7/3/7 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle ls-7/0/0.0;
        }
    }
}
[edit class-of-service]
classifiers {
    dscp dscp_default {
        import default;
    }
    inet-precedence inet-precedence_default {
        import default;
    }
```

```
}
code-point-aliases {
    dscp {
        af11 001010;
        af12 001100;
        af13 001110;
        af21 010010;
        af22 010100;
        af23 010110;
        af31 011010;
        af32 011100;
        af33 011110;
        af41 100010;
        af42 100100;
        af43 100110;
        be 000000;
        cs1 001000;
        cs2 010000;
        cs3 011000;
        cs4 100000;
        cs5 101000;
        cs6 110000;
        cs7 111000;
        ef 101110;
    }
    inet-precedence {
        af11 001;
        af21 010;
        af31 011;
        af41 100;
        be 000;
        cs6 110;
        cs7 111;
        ef 101;
        nc1 110;
        nc2 111;
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
```

```
}
interfaces {
    ge-7/2/0 {
        scheduler-map sched-map;
        unit 0 {
            classifiers {
                dscp dscp_default;
            }
        }
    }
    e1-7/3/6 {
        scheduler-map sched-map;
    }
    e1-7/3/7 {
        scheduler-map sched-map;
    }
    ls-7/0/0 {
        scheduler-map sched-map;
        unit 0 {
            classifiers {
                inet-precedence inet-precedence_default;
            }
        }
    }
}
scheduler-maps {
    sched-map {
        forwarding-class af scheduler af-scheduler;
        forwarding-class be scheduler be-scheduler;
        forwarding-class ef scheduler ef-scheduler;
        forwarding-class nc scheduler nc-scheduler;
    }
}
schedulers {
    af-scheduler {
        transmit-rate percent 25;
        buffer-size percent 25;
    }
    be-scheduler {
        transmit-rate percent 25;
        buffer-size percent 25;
    }
    ef-scheduler {
```

```
        transmit-rate percent 25;
        buffer-size percent 25;
    }
    nc-scheduler {
        transmit-rate percent 25;
        buffer-size percent 25;
    }
}
[edit firewall]
filter lfi_ls_filter {
    term term0 {
        from {
            destination-address {
                192.168.1.3/32;
            }
            precedence 5;
        }
        then {
            count count-192-168-1-3;
            forwarding-class af;
            accept;
        }
    }
    term default {
        then {
            log;
            forwarding-class best effort;
            accept;
        }
    }
}
```

RELATED DOCUMENTATION

Link and Multilink Services Overview | 2

Configuring Link Services Physical Interfaces | 134

Example: Configuring a Multilink Interface with MLPPP | 55

CHAPTER 4

Configuration Statements and Operational Commands

IN THIS CHAPTER

- Junos CLI Reference Overview | **149**

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also 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