Link Services IQ Interfaces

Published: 2013-02-15

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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

Table of Contents

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

For the features described in this document, the following platforms are supported:

- M Series
- T Series
- MX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xsl; }
```


2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

   [edit]
   user@host# **edit system scripts**
   [edit system scripts]

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

   [edit system scripts]
   user@host# **load merge relative /var/tmp/ex-script-snippet.conf**
   load complete

For more information about the **load** command, see the CLI User Guide.

Documentation Conventions

Table 1 on page xi defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page xi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the**configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`

Table 2: Text and Syntax Conventions *(continued)*

<table>
<tr><th>Convention</th><th>Description</th><th>Examples</th></tr>
<tr><td>Italic text like this</td><td>• Introduces or emphasizes important new terms.
• Identifies book names.
• Identifies RFC and Internet draft titles.</td><td>• A policy term is a named structure that defines match conditions and actions.
• Junos OS System Basics Configuration Guide
• RFC 1997, BGP Communities Attribute</td></tr>
<tr><td>Italic text like this</td><td>Represents variables (options for which you substitute a value) in commands or configuration statements.</td><td>Configure the machine's domain name:

[edit]
root@# set system domain-name domain-name</td></tr>
<tr><td>Text like this</td><td>Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.</td><td>• To configure a stub area, include the stub statement at the[edit protocols ospf area area-id] hierarchy level.
• The console port is labeled CONSOLE.</td></tr>
<tr><td>< > (angle brackets)</td><td>Enclose optional keywords or variables.</td><td>stub <default-metric <i>metric</i>>;</td></tr>
<tr><td>| (pipe symbol)</td><td>Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.</td><td>broadcast | multicast

(string1 | string2 | string3)</td></tr>
<tr><td># (pound sign)</td><td>Indicates a comment specified on the same line as the configuration statement to which it applies.</td><td>rsvp { # Required for dynamic MPLS only</td></tr>
<tr><td>[ ] (square brackets)</td><td>Enclose a variable for which you can substitute one or more values.</td><td>community name members [community-ids]</td></tr>
<tr><td>Indention and braces ({ })</td><td>Identify a level in the configuration hierarchy.</td><td rowspan="2">[edit]
routing-options {
  static {
    route default {
      nexthop address;
      retain;
    }
  }
}</td></tr>
<tr><td>; (semicolon)</td><td>Identifies a leaf statement at a configuration hierarchy level.</td></tr>
<tr><td colspan="3">J-Web GUI Conventions</td></tr>
<tr><td>Bold text like this</td><td>Represents J-Web graphical user interface (GUI) items you click or select.</td><td>• In the Logical Interfaces box, select All Interfaces.
• To cancel the configuration, click Cancel.</td></tr>
<tr><td>> (bold right angle bracket)</td><td>Separates levels in a hierarchy of J-Web selections.</td><td>In the configuration editor hierarchy, select Protocols>Ospf.</td></tr>
</table>

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

PART 1

Overview

CHAPTER 1

Link Services IQ interfaces

Layer 2 Service Package Capabilities and Interfaces

As described in Enabling Service Packages, you can configure the AS or Multiservices PIC and the internal ASM in the M7i platform to use either the Layer 2 or the Layer 3 service package.

When you enable the Layer 2 service package, the AS or Multiservices PIC supports *link services*. On the AS or Multiservices PIC and the ASM, link services include the following:

- Junos CoS components—"Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 19 describes how the Junos CoS components work on link services IQ (**lsq**) interfaces. For detailed information about Junos CoS components, see the Junos OS Class of Service Configuration Guide.
- Data compression using the compressed Real-Time Transport Protocol (CRTP) for use in voice over IP (VoIP) transmission.

NOTE: On LSQ interfaces, all multilink traffic for a single bundle is sent to a single processor. If CRTP is enabled on the bundle, it adds overhead to the CPU. Because T3 network interfaces support only one link per bundle, make sure you configure a fragmentation map for compressed traffic on these interfaces and specify the no-fragmentation option. For more information, see Configuring Delay-Sensitive Packet Interleaving and "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

- Link fragment interleaving (LFI) on Frame Relay links using FRF.12 end-to-end fragmentation—The standard for FRF.12 is defined in the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*.
- LFI on Multilink Point-to-Point Protocol (MLPPP) links.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—The standard for FRF.15 is defined in the specification FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*.

- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—The standard for FRF.16 is defined in the specification FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*.
- MLPPP—The standard for MLPPP is defined in the specification RFC 1990, *The PPP Multilink Protocol (MP)*.
- Multiclass extension to MLPPP—The standard is defined in the specification RFC 2686, *The Multi-Class Extension to Multi-Link PPP*.

For the LSQ interface on the AS or Multiservices PIC, the configuration syntax is almost the same as for Multilink and Link Services PICs. The primary difference is the use of the interface-type descriptor **lsq** instead of **ml** or **ls**. When you enable the Layer 2 service package on the AS or Multiservices PIC, the following interfaces are automatically created:

gr-*fpc/pic/port*
ip-*fpc/pic/port*
lsq-*fpc/pic/port*
lsq-*fpc/pic/port*:0
...
lsq-*fpc/pic/port*:*N*
mt-*fpc/pic/port*
pd-*fpc/pic/port*
pe-*fpc/pic/port*
sp-*fpc/pic/port*
vt-*fpc/pic/port*

Interface types **gr**, **ip**, **mt**, **pd**, **pe**, and **vt** are standard tunnel interfaces that are available on the AS or Multiservices PIC whether you enable the Layer 2 or the Layer 3 service package. These tunnel interfaces function the same way for both service packages, except that the Layer 2 service package does not support some tunnel functions, as shown in Table 5 on page 24. For more information about tunnel interfaces, see Tunnel Properties.

NOTE: Interface type **sp** is created because it is needed by the Junos OS. For the Layer 2 service package, the **sp** interface is not configurable, but you should not disable it.

Interface type **lsq-*fpc/pic/port*** is the physical link services IQ interface (**lsq**). Interface types **lsq-*fpc/pic/port*:0** through **lsq-*fpc/pic/port*:*N*** represent FRF.16 bundles. These interface types are created when you include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level. For more information, see "Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 19.

NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the router does not use the bandwidth value if the interfaces are included in an MLPPP or MLFR bundle. The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about these properties, see the Junos® OS Network Interfaces.

PART 2

Configuration

CHAPTER 2

Configuration Tasks

Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS

Link services IQ (**lsq-**) interfaces that are paired with SONET PICs can use the Automatic Protection Switching (APS) configuration already available on SONET networks to provide failure recovery. SONET APS provides stateless failure recovery, if it is configured on SONET interfaces in separate chassis and each SONET PIC is paired with an AS or Multiservices PIC in the same chassis. If one of the following conditions for APS failure is met, the associated SONET PIC triggers recovery to the backup circuit and its associated AS or Multiservices PIC. The failure conditions are:

- Failure of Link Services IQ PIC
- Failure of FPC that hosts the Link Services IQ PIC
- Failure of Packet Forwarding Engine
- Failure of chassis

The guidelines for configuring SONET APS are described in the Junos® OS Network Interfaces.

The following sections describe how to configure failover properties:

Configuring the Association between LSQ and SONET Interfaces

To configure the association between AS or Multiservices PICs hosting link services IQ interfaces and the SONET interfaces, include the **lsq-failure-options** statement at the **[edit interfaces]** hierarchy level:

```
lsq-fpc/pic/port {
  lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
  }
}
```

For example, consider the following network scenario:

- Primary router includes interfaces **oc3-0/2/0** and **lsq-1/1/0**.
- Backup router includes interfaces **oc3-2/2/0** and **lsq-3/2/0**.

Configure SONET APS, with **oc3-0/2/0** as the working circuit and **oc3-2/2/0** as the protect circuit. Include the **trigger-link-failure** statement to extend failure to the LSQ PICs:

```
interfaces lsq-1/1/0 {
  lsq-failure-options {
    trigger-link-failure oc3-0/2/0;
  }
}
```

NOTE: You must configure the **lsq-failure-options** statement on the primary router only. The configuration is not supported on the backup router.

To inhibit the router from sending PPP termination-request messages to the remote host if the Link Services IQ PIC fails, include the **no-termination-request** statement at the **[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]** hierarchy level:

[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]
no-termination-request;

This functionality is supported on link PICs as well. To inhibit the router from sending PPP termination-request messages to the remote host if a link PIC fails, include the **no-termination-request** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

[edit interfaces *interface-name* ppp-options]
no-termination-request;

The **no-termination-request** statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only, on the following PICs:

- Channelized OC3 IQ PICs
- Channelized OC12 IQ PICs
- Channelized STM1 IQ PICs
- Channelized STM4 IQ PICs

Configuring SONET APS Interoperability with Cisco Systems FRF.16

Juniper Networks routers configured with APS might not interoperate correctly with Cisco FRF.16. To enable interoperation, include the **cisco-interoperability** statement at the **[edit interfaces lsq-*fpc/pic/port* mlfr-uni-nni-bundle-options]** hierarchy level:

[edit interfaces lsq-*fpc/pic/port* mlfr-uni-nni-bundle-options]
cisco-interoperability send-lip-remove-link-for-link-reject;

The **send-lip-remove-link-for-link-reject** option prompts the router to send a Link Integrity Protocol remove link when it receives an add-link rejection message.

Restrictions on APS Redundancy for LSQ Interfaces

The following restrictions apply to LSQ failure recovery:

- It applies only to Link Services IQ PICs installed in M Series routers, except for M320 routers.
- You must configure the **failure-options** statement on physical LSQ interfaces, not on MLFR channelized units.
- The Link Services IQ PICs must be associated with SONET link PICs. The paired PICs can be installed on different routers or in the same router; in other words, both interchassis and intrachassis recovery are supported
- Failure recovery is stateless; as a result, route flapping and loss of link state is expected in interchassis recovery, requiring PPP renegotiation. In intrachassis recovery, no impact on traffic is anticipated with Routing Engine failover, but PIC failover results in PPP renegotiation.

- The switchover is not revertive: when the original hardware is restored to service, traffic does not automatically revert back to it.
- Normal APS switchover and PIC-triggered APS switchover can be distinguished only by checking the system log messages.

NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

Configuring LSQ Interface Redundancy in a Single Router Using SONET APS

Stateless switchover from one Link Services IQ PIC to another within the same router can be configured by using the SONET APS mechanism described in "Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS" on page 7. Each Link Services IQ PIC must be associated with a specified SONET link PIC within the same router.

NOTE: For complete intrachassis recovery, including recovery from Routing Engine failover, graceful Routing Engine switchover (GRES) must be enabled on the router. For more information, see the Junos OS System Basics Configuration Guide.

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

You can configure failure recovery on M Series, MX Series, and T Series routers that have multiple AS or Multiservices PICs and DPCs with **lsq-** interfaces by specifying a virtual LSQ redundancy (**rlsq**) interface in which the primary Link Services IQ PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all LSQ processing is transferred to it. To determine which PIC is currently active, issue the **show interfaces redundancy** command.

NOTE: This configuration does not require the use of SONET APS for failover. Network interfaces that do not support SONET can be used, such as T1 or E1 interfaces.

The following sections provide more information:

- Configuring Redundant Paired LSQ Interfaces on page 11
- Restrictions on Redundant LSQ Interfaces on page 12
- Configuring Link State Replication for Redundant Link PICs on page 13
- Examples: Configuring Redundant LSQ Interfaces for Failure Recovery on page 15

Configuring Redundant Paired LSQ Interfaces

The physical interface type **rlsq** specifies the pairings between primary and secondary **lsq** interfaces to enable redundancy. To configure a backup **lsq** interface, include the **redundancy-options** statement at the **[edit interfaces rlsq*number*]** hierarchy level:

```
[edit interfaces rlsqnumber]
redundancy-options {
  (hot-standby | warm-standby);
  primary lsq-fpc/pic/port;
  secondary lsq-fpc/pic/port;
}
```

For the **rlsq** interface, ***number*** can be from 0 through 1023. If the primary **lsq** interface fails, traffic processing switches to the secondary interface. The secondary interface remains active even after the primary interface recovers. If the secondary interface fails and the primary interface is active, processing switches to the primary interface.

The **hot-standby** option is used with one-to-one redundancy configurations, in which one working PIC is supported by one backup PIC. It is supported with MLPPP, CRTP, FRF.15, and FRF.16 configurations for the LSQ interface to achieve an uninterrupted LSQ service. It sets the requirement for the failure detection and recovery time to be less than 5 seconds. The behavior is revertive, but you can manually switch between the primary and secondary PICs by issuing the **request interfaces (revert | switchover) rlsq*number*** operational mode command. It also provides a switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

The **warm-standby** option is used with redundancy configurations in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected.

Certain combinations of **hot-standby** and **warm-standby** configuration are not permitted and result in a configuration error. The following examples are permitted:

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **warm-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq0:1** configured with **primary lsq-0/0/0:1**

The following example combinations are not permitted:

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **hot-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq1:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:1**, in combination with interface **rlsq1:1** configured with **primary lsq-0/0/0:1**
- Interface **rlsq0** configured with **primary lsq-0/0/0**, in combination with interface **rlsq1** configured with **primary lsq-0/0/0**

In addition, the same physical interface cannot be reused as the primary interface for more than one **rlsq** interface, nor can any of the associated logical interfaces. For example, primary interface **lsq-0/0/0** cannot be reused in another **rlsq** interface as **lsq-0/0/0:0**.

Restrictions on Redundant LSQ Interfaces

Link Services IQ PIC failure occurs under the following conditions:

- The primary PIC fails to boot. In this case, the **rlsq** interface does not come up and manual intervention is necessary to reboot or replace the PIC, or to rename the primary PIC to the secondary one in the **rlsq** configuration.
- When configuring an **rlsq** interface, ensure that:
  - The unit number allocated to the **rlsq** interface is less than the number of Multilink Frame Relay user-to-network interface network-to-network interface (UNI-NNI) (FRF.16) bundles allocated on the Link Services PIC.
  - Data-link connection identifier (DLCI) is configured for the **rlsq** interface.

  If these conditions are not met, the **rlsq** interface does not boot. When you issue the **show interfaces redundancy** command, the state of the **rlsq** interface is indicated as **Waiting for primary MS PIC**.
- The primary PIC becomes active and then fails. The secondary PIC automatically takes over processing.
- A failover to the secondary PIC takes place. The secondary PIC then fails. If the primary PIC has been restored to active state, processing switches to it.
- The FPC that contains the Link Services IQ PIC fails.

The following constraints apply to redundant LSQ configurations:

- We recommend that primary and secondary PICs be configured in two different FPCs (in chassis other than M10i routers).
- You cannot configure a Link Services IQ PIC with explicit bundle configurations and as a constituent of an **rlsq** interface.
- Redundant LSQ configurations provide full GRES support. (You must configure GRES at the **[edit chassis]** hierarchy level; see the Junos OS System Basics Configuration Guide.
- If you configure the **redundancy-options** statement with the **hot-standby** option, the configuration must include one **primary** interface value and one **secondary** interface value.
- Since the same interface name is used for **hot-standby** and **warm-standby**, if you modify the configuration to change this attribute, it is recommended that you first deactivate the interface, commit the new configuration, and then reactivate the interface.
- You cannot make changes to an active **redundancy-options** configuration. You must deactivate the **rlsq*number*** interface configuration, change it, and reactivate it.

- The **rlsq*number*** configuration becomes active only if the primary interface is active. When the configuration is first activated, the primary interface must be active; if not, the **rlsq** interface waits until the primary interface comes up.
- You cannot modify the configuration of **lsq** interfaces after they have been included in an active **rlsq** interface.
- All the operational mode commands that apply to **rsp** interfaces also apply to **rlsq** interfaces. You can issue **show** commands for the **rlsq** interface or the primary and secondary **lsq** interfaces. However, statistics on the link interfaces are not carried over following a Routing Engine switchover.
- The **rlsq** interfaces also support the **lsq-failure-options** configuration, discussed in "Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS" on page 7. If the primary and secondary Link Services IQ PICs fail and the **lsq-failure-options** statement is configured, the configuration triggers a SONET APS switchover.
- Redundant LSQ configurations that require MLPPP Multilink Frame Relay (FRF.15 and FRF.16) are supported only with the **warm-standby** option.
- Redundant LSQ support is extended to ATM network interfaces.
- Channelized interfaces are used with FRF-16 bundles, for example **rlsq0:0**. The **rlsq** number and its constituents, the **primary** and **secondary** interfaces, must match for the configuration to be valid: either all must be channelized, or none. For an example of an FRF.16 configuration, see "Configuring LSQ Interface Redundancy for an FRF.16 Bundle" on page 18.

NOTE: **Adaptive Services and Multiservices PICs in layer-2 mode (running Layer 2 services) are not rebooted when a MAC flow-control situation is detected.**

Configuring Link State Replication for Redundant Link PICs

Link state replication, also called *interface preservation*, is an addition to the SONET Automatic Protection Switching (APS) functionality that helps promote redundancy of the link PICs used in LSQ configurations.

Link state replication provides the ability to add two sets of links, one from the active (working) SONET PIC and the other from the backup (protect) SONET PIC to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without causing a link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation. For more information about SONET APS configurations, see the Junos® OS Network Interfaces.

To configure link state replication, include the **preserve-interface** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level on both network interfaces:

```
edit interfaces interface-name sonet-options aps]
preserve-interface;
```

The following constraints apply to link PIC redundancy:

- APS functionality must be available on the SONET PICs and the interface configurations must be identical on both ends of the link. Any configuration mismatch causes the commit operation to fail.
- This feature is supported only with LSQ and SONET APS-enabled link PICs, including Channelized OC3, Channelized OC12, and Channelized STM1 intelligent queuing (IQ) PICs.
- Link state replication supports MLPPP and PPP over Frame Relay (**frame-relay-ppp**) encapsulation, and fully supports GRES.
- Enabling the interface or protocol traceoptions with a large number of MLPPP links can trigger Link Control Protocol (LCP) renegotiation during the link switchover time.

  NOTE: This renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an add/drop multiplexer (ADM).

- In general, networks that connect a Juniper Networks router to an ADM allow faster MLPPP link switchover than those with back-to-back Juniper Networks routers. The MLPPP link switchover time difference may be significant, especially for networks with a large number of MLPPP links.
- An aggressive LCP keepalive timeout configuration can lead to LCP renegotiation during the MLPPP link switchover. By default, the LCP keepalive timer interval is 10 seconds and the consecutive link down count is 3. The MLPPP links start LCP negotiation only after a timeout of 30 seconds. Lowering these configuration values may trigger one or more of the MLPPP links to renegotiate during the switchover time.

  NOTE: LCP renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an ADM.

As an example, the following configuration shows the link state replication configuration between the ports **coc3-1/0/0** and c**oc3-2/0/0**.

```
interfaces {
  coc3-1/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        working-circuit aps-group-1;
      }
    }
  }
  coc3-2/0/0 {
    sonet-options {
      aps {
        preserve-interface;
```

```
                protect-circuit aps-group-1;
            }
        }
    }
}
```

Examples: Configuring Redundant LSQ Interfaces for Failure Recovery

Configuring LSQ Interface Redundancy for MLPPP

The following configuration shows that **lsq-1/1/0** and **lsq-1/3/0** work as a pair and the redundancy type is **hot-standby**, which sets the requirement for the failure detection and recovery time to be less than 5 seconds:

```
interfaces rlsq0 {
    redundancy-options {
        primary lsq-1/1/0;
        secondary lsq-1/3/0;
        hot-standby; #either hot-standby or warm-standby is supported
    }
}
```

The following example shows a related MLPPP configuration:

NOTE: MLPPP protocol configuration is required for this configuration.

```
interfaces {
    t1-/1/2/0 {
        unit 0 {
            family mlppp {
                bundle rlsq0.0;
            }
        }
    }
    rlsq0 {
        unit 0 {
            family inet {
                address 30.1.1.2/24;
            }
        }
    }
}
```

The following example shows a related CoS configuration:

```
class-of-service {
    interfaces {
        rlsq0 {
            unit * {
                fragmentation-maps fr-map1;
            }
        }
    }
```

```
}
```

The following example shows a complete link state replication configuration for MLPPP. This example uses two bundles, each with four T1 links. The first four T1 links (**t1-*:1** through **t1-*:4**) form the first bundle and the last four T1 links (**t1-*:5** through **t1-*:8**) form the second bundle. To minimize the duplication in the configuration, this example uses the **[edit groups]** statement; for more information, see the Junos OS System Basics Configuration Guide. This type of configuration is not required; it simplifies the task and minimizes duplication.

```
groups {
   ml-partition-group {
      interfaces {
         <coc3-*> {
            partition 1 oc-slice 1 interface-type coc1;
         }
         <coc1-*> {
            partition 1-8 interface-type t1;
         }
      }
   }
   ml-bundle-group-1 {
      interfaces {
         <t1-*:"[1-4]"> {
            encapsulation ppp;
            unit 0 {
               family mlppp {
                  bundle lsq-0/1/0.0;
               }
            }
         }
      }
   }
   ml-bundle-group-2 {
      interfaces {
         <t1-*:"[5-8]"> {
            encapsulation ppp;
            unit 0 {
               family mlppp {
                  bundle lsq-0/1/0.1;
               }
            }
         }
      }
   }
}
interfaces {
   lsq-0/1/0 {
      unit 0 {
         encapsulation multilink-ppp;
         family inet {
            address 1.1.1.1/32 {
               destination 1.1.1.2;
            }
         }
```

```
    }
    unit 1 {
        encapsulation multilink-ppp;
        family inet {
            address 1.1.2.1/32 {
                destination 1.1.2.2;
            }
        }
    }
}
coc3-1/0/0 {
    apply-groups ml-partition-group;
    sonet-options {
        aps {
            preserve-interface;
            working-circuit aps-group-1;
        }
    }
}
coc1-1/0/0:1 {
    apply-groups ml-partition-group;
}
t1-1/0/0:1:1 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:2 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:3 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:4 {
    apply-groups ml-bundle-group-1;
}
t1-1/0/0:1:5 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:6 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:7 {
    apply-groups ml-bundle-group-2;
}
t1-1/0/0:1:8 {
    apply-groups ml-bundle-group-2;
}
coc3-2/0/0 {
    apply-groups ml-partition-group;
    sonet-options {
        aps {
            preserve-interface;
            protect-circuit aps-group-1;
        }
    }
}
coc1-2/0/0:1 {
```



```
      apply-groups ml-partition-group;
    }
    t1-2/0/0:1:1 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:2 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:3 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:4 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:5 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:6 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:7 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:8 {
      apply-groups ml-bundle-group-2;
    }
  }
```

Configuring LSQ Interface Redundancy for an FRF.15 Bundle

The following example shows a configuration for an FRF.15 bundle:

```
interfaces rlsq0 {
  redundancy-options {
    primary lsq-1/2/0;
    secondary lsq-1/3/0;
    warm-standby; #either hot-standby or warm-standby is supported
  }
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 30.1.1.1/24;
    }
  }
}
```

Configuring LSQ Interface Redundancy for an FRF.16 Bundle

The following example shows a configuration for an FRF.16 bundle:

```
interfaces rlsq0:0 {
  dce;
  encapsulation multilink-frame-relay-uni-nni;
```

```
    redundancy-options {
      primary lsq-1/2/0:0;
      secondary lsq-1/3/0:0;
      warm-standby; #either hot-standby or warm-standby is supported
    }
    unit 0 {
      dlci 1000;
      family inet {
        address 50.1.1.1/24;
      }
    }
  }
```

Configuring CoS Scheduling Queues on Logical LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify a scheduler map for each logical unit. A logical unit represents either an MLPPP bundle or a DLCI configured on a FRF.16 bundle. The scheduler is applied to the traffic sent to an AS or Multiservices PIC running the Layer 2 link services package.

If you configure a scheduler map on a bundle, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*]** hierarchy level. If you configure a scheduler map on an FRF.16 DLCI, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*:*channel*]** hierarchy level. For more information, see the Junos OS Class of Service Configuration Guide.

If you need latency guarantees for multiclass or LFI traffic, you must use channelized IQ PICs for the constituent links. With non-IQ PICs, because queueing is not done at the channelized interface level on the constituent links, latency-sensitive traffic might not receive the type of service that it should. Constituent links from the following PICs support latency guarantees:

- Channelized E1 IQ PIC
- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC
- Channelized T3 IQ PIC

For scheduling queues on a logical interface, you can configure the following scheduler map properties at the **[edit class-of-service schedulers]** hierarchy level:

- **buffer-size**—The queue size; for more information, see "Configuring Scheduler Buffer Size" on page 20.
- **priority**—The transmit priority (low, high, strict-high); for more information, see "Configuring Scheduler Priority" on page 21.

- **shaping-rate**—The subscribed transmit rate; for more information, see "Configuring Scheduler Shaping Rate" on page 21.
- **drop-profile-map**—The random early detection (RED) drop profile; for more information, see "Configuring Drop Profiles" on page 21.

When you configure MLPPP and FRF.12 on M Series and T Series routers, you should configure a single scheduler with non-zero 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.

When you configure FRF.16 on M Series and T Series routers, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16" on page 46. For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bandwidth to be used by queue 0. The default scheduler transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent, respectively. This default scheduler sends all user traffic to queue 0 and all network-control traffic to queue 3, and therefore it is well suited to the behavior of FRF.16. You can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behaviors, and apply it to the constituent links.

NOTE: **On T Series and M320 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.**

For link services IQ interfaces (**lsq**), these scheduling properties work as they do in other PICs, except as noted in the following sections.

NOTE: **On T Series and M320 routers, lsq interfaces do not support DiffServ code point (DSCP) and DSCP-IPv6 rewrite markers.**

Configuring Scheduler Buffer Size

You can configure the scheduler buffer size in three ways: as a temporal value, as a percentage, and as a remainder. On a single logical interface (MLPPP or a FRF.16 DLCI), each queue can have a different buffer size.

If you specify a temporal value, the queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This number is computed by multiplying logical interface speed by the temporal value. For MLPPP bundles, logical interface speed is equal to the bundle bandwidth, which is the sum of constituent link speeds minus link-layer overhead. For MLFR FRF.16 DLCIs, logical interface speed is equal to bundle bandwidth multiplied by the DLCI shaping rate. In all cases, the maximum temporal value is limited to 200 milliseconds.

Buffer size percentages are implicitly converted into temporal values by multiplying the percentage by 200 milliseconds. For example, buffer size specified as **buffer-size percent 20** is the same as a 40-millisecond temporal delay. The link services IQ implementation guarantees 200 milliseconds of buffer delay for all interfaces with T1 and higher speeds. For slower interfaces, it guarantees one second of buffer delay.

The queueing algorithm evenly distributes leftover bandwidth among all queues that are configured with the **buffer-size remainder** statement. The queuing algorithm guarantees enough space in the transmit buffer for two MTU-sized packets.

Configuring Scheduler Priority

The transmit priority of each queue is determined by the scheduler and the forwarding class. Each queue receives a guaranteed amount of bandwidth specified with the scheduler **transmit-rate** statement.

Configuring Scheduler Shaping Rate

You use the shaping rate to set the percentage of total bundle bandwidth that is dedicated to a DLCI. For link services IQ DLCIs, only percentages are accepted, which allows adjustments in response to dynamic changes in bundle bandwidth—for example, when a link goes up or down. This means that absolute shaping rates are not supported on FRF.16 bundles. Absolute shaping rates are allowed for MLPPP and MLFR bundles only.

For scheduling between DLCIs in a MLFR FRF.16 bundle, you can configure a shaping rate for each DLCI. A shaping rate is expressed as a percentage of the aggregate bundle bandwidth. Shaping rate percentages for all DLCIs within a bundle can add up to 100 percent or less. Leftover bandwidth is distributed equally to DLCIs that do not have the **shaping-rate** statement included at the **[edit class-of-service interfaces lsq-*fpc/pic/port*:*channel* unit *logical-unit-number*]** hierarchy level. If none of the DLCIs in an MLFR FRF.16 bundle specify a DLCI scheduler, the total bandwidth is evenly divided across all DLCIs.

NOTE: **For FRF.16 bundles on link services IQ interfaces, only shaping rates based on percentage are supported.**

Configuring Drop Profiles

You can configure random early detection (RED) on LSQ interfaces as in other CoS scenarios. To configure RED, include one or more drop profiles and attach them to a scheduler for a particular forwarding class. For more information about RED profiles, see the Junos OS Class of Service Configuration Guide.

The LSQ implementation performs tail RED. It supports a maximum of 256 drop profiles per PIC. Drop profiles are configurable on a per-queue, per-loss-priority, and per-TCP-bit basis.

You can attach scheduler maps with configured RED drop profiles to any LSQ logical interface: an MLPPP bundle, an FRF.15 bundle, or an FRF.16 DLCI. Different queues (forwarding classes) on the same logical interface can have different associated drop profiles.

The following example shows how to configure a RED profile on an LSQ interface:

```
[edit]
class-of-service {
  drop-profiles {
    drop-low {
      # Configure suitable drop profile for low loss priority
      ...
    }
    drop-high {
      # Configure suitable drop profile for high loss priority
      ...
    }
  }
  scheduler-maps {
    schedmap {
      # Best-effort queue will use be-scheduler
      # Other queues may use different schedulers
      forwarding-class be scheduler be-scheduler;
      ...
    }
  }
  schedulers {
    be-scheduler {
      # Configure two drop profiles for low and high loss priority
      drop-profile-map loss-priority low protocol any drop-profile drop-low;
      drop-profile-map loss-priority high protocol any drop-profile drop-high;
      # Other scheduler parameters (buffer-size, priority,
      # and transmit-rate) are already supported.
      ...
    }
  }
  interfaces {
    lsq-1/3/0.0 {
      # Attach a scheduler map (that includes RED drop profiles)
      # to a LSQ logical interface.
      scheduler-map schedmap;
    }
  }
}
```

NOTE: The RED profiles should be applied only on the LSQ bundles and not on the egress links that constitute the bundle.

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify fragmentation properties for specific forwarding classes. Traffic on each forwarding class can be either multilink encapsulated (fragmented and sequenced) or nonencapsulated (hashed with no fragmentation). By default, traffic in all forwarding classes is multilink encapsulated.

When you do not configure fragmentation properties for the queues on MLPPP interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number* fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLPPP interface. For MLFR FRF.16 interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLFR FRF.16 interface.

If you do not set a maximum fragment size anywhere in the configuration, packets are still fragmented if they exceed the smallest maximum transmission unit (MTU) or maximum received reconstructed unit (MRRU) of all the links in the bundle. A nonencapsulated flow uses only one link. If the flow exceeds a single link, then the forwarding class must be multilink encapsulated, unless the packet size exceeds the MTU/MRRU.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the MRRU by including the **mrru** statement at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** 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.

To configure fragmentation properties on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      (fragment-threshold bytes | no-fragmentation);
      multilink-class number;
    }
  }
}
```

To set a per-forwarding class fragmentation threshold, include the **fragment-threshold** statement in the fragmentation map. This statement sets the maximum size of each multilink fragment.

To set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. This statement specifies that an extra fragmentation header is not prepended to the packets received on this queue and that static link load balancing is used to ensure in-order packet delivery.

For a given forwarding class, you can include either the **fragment-threshold** or **no-fragmentation** statement; they are mutually exclusive.

You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For a given forwarding class, you can include either the **multilink-class** or

no-fragmentation statement; they are mutually exclusive. For more information about MCML, see "Configuring Multiclass MLPPP on LSQ Interfaces" on page 25.

To associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI, include the **fragmentation-map** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces]
lsq-fpc/pic/port {
  unit logical-unit-number { # Multilink PPP
  fragmentation-map map-name;
}
lsq-fpc/pic/port:channel { # MLFR FRF.16
unit logical-unit-number {
  fragmentation-map map-name;
}
```

For configuration examples, see the following topics:

- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 37
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 43
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 48
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 53
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 60
- Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 61
- Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 62
- Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 64

For Link Services PIC link services (**ls-**) interfaces, fragmentation maps are not supported. Instead, you enable LFI by including the **interleave-fragments** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. For more information, see Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces.

Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces

Link-layer overhead can cause packet drops on constituent links because of bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information.

By default, 4 percent of the total bundle bandwidth is set aside for link-layer overhead. In most network environments, the average link-layer overhead is 1.6 percent. Therefore, we recommend 4 percent as a safeguard. For more information, see RFC 4814, *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking*.

For link services IQ (**lsq-**) interfaces, you can configure the percentage of bundle bandwidth to be set aside for link-layer overhead. To do this, include the **link-layer-overhead** statement:

link-layer-overhead *percent*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]**
- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You can configure the value to be from 0 percent through 50 percent.

Configuring Multiclass MLPPP on LSQ Interfaces

For link services IQ (**lsq-**) interfaces with MLPPP encapsulation, you can configure multiclass MLPPP (MCML). If you do not configure MCML, 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 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.

For link services IQ interfaces only, you can configure MCML, as defined in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. MCML makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, MCML allows different classes of traffic to have different latency guarantees. With MCML, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees.

NOTE: Configuring both LFI and MCML 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.

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

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

To configure MCML on a link services IQ 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 MCML class.

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.

To specify the mapping of a forwarding class into a MCML 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.

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

Oversubscribing Interface Bandwidth on LSQ Interfaces

The term *oversubscribing interface bandwidth* means configuring shaping rates (peak information rates [PIRs]) so that their sum exceeds the interface bandwidth.

On Channelized IQ PICs, Gigabit Ethernet IQ PICs, and FRF.16 link services IQ (**lsq-**) interfaces on AS and Multiservices PICs, you can oversubscribe interface bandwidth. The logical interfaces (and DLCIs within an FRF.16 bundle) can be oversubscribed when there is leftover bandwidth. The oversubscription is limited to the configured PIR. Any unused bandwidth is distributed equally among oversubscribed logical interfaces or DLCIs.

For networks that are not likely to experience congestion, oversubscribing interface bandwidth improves network utilization, thereby allowing more customers to be provisioned on a single interface. If the actual data traffic does not exceed the interface bandwidth, oversubscription allows you to sell more bandwidth than the interface can support.

We recommend avoiding oversubscription in networks that are likely to experience congestion. Be careful not to oversubscribe a service by too much, because this can cause degradation in the performance of the router during congestion. When you configure oversubscription, some output queues can be starved if the actual data traffic exceeds the physical interface bandwidth. You can prevent degradation by using statistical multiplexing to ensure that the actual data traffic does not exceed the interface bandwidth.

NOTE: You cannot oversubscribe interface bandwidth when you configure traffic shaping using the method described in Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs.

When configuring oversubscription for FRF.16 bundle interfaces, you can assign traffic control profiles that apply on a physical interface basis. When you apply traffic control profiles to FRF.16 bundles at the *logical* interface level, member link interface bandwidth is underutilized when there is a small proportion of traffic or no traffic at all on an individual DLCI. Support for traffic control features on the FRF.16 bundle physical interface level addresses this limitation.

To configure oversubscription of an interface, perform the following steps:

1. Include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   shaping-rate (percent *percentage* | *rate*);

   NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **shaping-rate** as a percentage.

   On LSQ interfaces, you can configure the shaping rate as a percentage.

   On IQ and IQ2 interfaces, you can configure the shaping rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

   Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. However, with this configuration approach, you cannot independently control the delay-buffer rate, as described in Step 2.

NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see "Configuring Guaranteed Minimum Rate on LSQ Interfaces" on page 31.

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **delay-buffer-rate** as a percentage.

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   delay-buffer-rate (percent percentage | rate);
   ```

   The delay-buffer rate overrides the shaping rate as the basis for the delay-buffer calculation. In other words, the shaping rate or scaled shaping rate is used for delay-buffer calculations only when the delay-buffer rate is not configured.

   For LSQ interfaces, if you do not configure a delay-buffer rate, the guaranteed rate (CIR) is used to assign buffers. If you do not configure a guaranteed rate, the shaping rate (PIR) is used in the undersubscribed case, and the scaled shaping rate is used in the oversubscribed case.

   On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

   On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

   The actual delay buffer is based on the calculations described in the Junos OS Class of Service Configuration Guide. For an example showing how the delay-buffer rates are applied, see "Examples: Oversubscribing an LSQ Interface" on page 30.

   Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed.

   This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of zero, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If you do not configure a delay-buffer rate or a guaranteed rate, the logical interface receives a delay-buffer rate in proportion to the shaping rate and the remaining delay-buffer rate available. In other words, the delay-buffer rate for each logical interface with no configured delay-buffer rate is equal to:

(remaining delay-buffer rate * shaping rate) / (sum of shaping rates)

The remaining delay-buffer rate is equal to:

(interface speed) – (sum of configured delay-buffer rates)

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   scheduler-map *map-name*;

   For information about configuring schedulers and scheduler maps, see the Junos OS Class of Service Configuration Guide.

4. Optionally, you can enable large buffer sizes to be configured. To do this, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

   [edit chassis fpc *slot-number* pic *pic-number*]
   q-pic-large-buffer;

   If you do not include this statement, the delay-buffer size is more restricted. We recommend restricted buffers for delay-sensitive traffic, such as voice traffic. For more information, see the Junos OS Class of Service Configuration Guide.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   per-unit-scheduler;

   When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 384.

6. To enable scheduling for FRF.16 bundles physical interfaces, include the **no-per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   no-per-unit-scheduler;

7. To apply the traffic-scheduling profile to the logical interface, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]
output-traffic-control-profile *profile-name*;

You cannot include the **output-traffic-control-profile** statement in the configuration if any of the following statements are included in the logical interface configuration: **scheduler-map**, **shaping-rate**, **adaptive-shaper**, or **virtual-channel-group**.

For a table that shows how the bandwidth and delay buffer are allocated in various configurations, see the Junos OS Class of Service Configuration Guide.

Examples: Oversubscribing an LSQ Interface

Oversubscribing an LSQ Interface with Scheduling Based on the Logical Interface

Apply a traffic-control profile to a logical interface representing a DLCI on an FRF.16 bundle.

```
interfaces {
  lsq-1/3/0:0 {
    per-unit-scheduler;
    unit 0 {
      dlci 100;
    }
    unit 1 {
      dlci 200;
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tc_0 {
      shaping-rate percent 100;
      guaranteed-rate percent 60;
      delay-buffer-rate percent 80;
    }
    tc_1 {
      shaping-rate percent 80;
      guaranteed-rate percent 40;
    }
  }
  interfaces {
    lsq-1/3/0 {
      unit 0 {
        output-traffic-control-profile tc_0;
      }
      unit 1 {
        output-traffic-control-profile tc_1;
      }
    }
  }
}
```

Oversubscribing an LSQ Interface with Scheduling Based on the Physical Interface

Apply a traffic-control profile to the physical interface representing an FRF.16 bundle:

```
interfaces {
  lsq-0/2/0:0 {
    no-per-unit-scheduler;
    encapsulation multilink-frame-relay-uni-nni;
```

```
        unit 0 {
            dlci 100;
            family inet {
                address 18.18.18.2/24;
            }
        }
    }
class-of-service {
    traffic-control-profiles {
        rlsq_tc {
            scheduler-map rlsq;
            shaping-rate percent 60;
            delay-buffer-rate percent 10;
        }
    }
    interfaces {
        lsq-0/2/0:0 {
            output-traffic-control-profile rlsq_tc;
        }
    }
}
scheduler-maps {
    rlsq {
        forwarding-class best-effort scheduler rlsq_scheduler;
        forwarding-class expedited-forwarding scheduler rlsq_scheduler1;
    }
}
schedulers {
    rlsq_scheduler {
        transmit-rate percent 20;
        priority low;
    }
    rlsq_scheduler1 {
        transmit-rate percent 40;
        priority high;
    }
}
```

Configuring Guaranteed Minimum Rate on LSQ Interfaces

On Gigabit Ethernet IQ PICs, Channelized IQ PICs, and FRF.16 link services IQ (LSQ) interfaces on AS and Multiservices PICs, you can configure guaranteed bandwidth, also known as a committed information rate (CIR). This allows you to specify a guaranteed rate for each logical interface. The guaranteed rate is a minimum. If excess physical interface bandwidth is available for use, the logical interface receives more than the guaranteed rate provisioned for the interface.

You cannot provision the sum of the guaranteed rates to be more than the physical interface bandwidth, or the bundle bandwidth for LSQ interfaces. If the sum of the guaranteed rates exceeds the interface or bundle bandwidth, the commit operation does not fail, but the software automatically decreases the rates so that the sum of the guaranteed rates is equal to the available bundle bandwidth.

To configure a guaranteed minimum rate, perform the following steps:

1. Include the **guaranteed-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   guaranteed-rate (percent *percentage* | *rate*);

   On LSQ interfaces, you can configure the guaranteed rate as a percentage.

   On IQ and IQ2 interfaces, you can configure the guaranteed rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

   NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

   This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see the Junos OS Class of Service Configuration Guide.

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   delay-buffer-rate (percent *percentage* | *rate*);

   On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

   On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

   The actual delay buffer is based on the calculations described in tables in the Junos OS Class of Service Configuration Guide. For an example showing how the delay-buffer rates are applied, see "Example: Configuring Guaranteed Minimum Rate" on page 34.

   If you do not include the **delay-buffer-rate** statement, the delay-buffer calculation is based on the guaranteed rate, the shaping rate if no guaranteed rate is configured, or the scaled shaping rate if the interface is oversubscribed.

   If you do not specify a shaping rate or a guaranteed rate, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 4 MTU-sized packets.

   You can configure a rate for the delay buffer that is higher than the guaranteed rate. This can be useful when the traffic flow might not require much bandwidth in general, but in some cases can be bursty and therefore needs a large buffer.

Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed. This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of 0, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If the guaranteed rate of a logical interface cannot be implemented, that logical interface receives a delay-buffer rate of 0, even if the configured delay-buffer rate is within the interface speed. If at a later time the guaranteed rate of the logical interface can be met, the configured delay-buffer rate is reevaluated and if the delay-buffer rate is within the remaining bandwidth, it is implemented.

If any logical interface has a configured guaranteed rate, all other logical interfaces on that port that do not have a guaranteed rate configured receive a delay-buffer rate of 0. This is because the absence of a guaranteed rate configuration corresponds to a guaranteed rate of 0 and, consequently, a delay-buffer rate of 0.

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   scheduler-map *map-name*;

   For information about configuring schedulers and scheduler maps, see the Junos OS Class of Service Configuration Guide.

4. To enable large buffer sizes to be configured, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

   [edit chassis fpc *slot-number* pic *pic-number*]
   q-pic-large-buffer;

   If you do not include this statement, the delay-buffer size is more restricted. For more information, see the Junos OS Class of Service Configuration Guide.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   per-unit-scheduler;

   When you include this statement, the maximum number of VLANs supported is 767 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 383.

6. To apply the traffic-scheduling profile to the logical interface, include the output-traffic-control-profile statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]
   output-traffic-control-profile *profile-name*;

Example: Configuring Guaranteed Minimum Rate

Two logical interface units, **0** and **1**, are provisioned with a guaranteed minimum of 750 Kbps and 500 Kbps, respectively. For logical unit **1**, the delay buffer is based on the guaranteed rate setting. For logical unit **0**, a delay-buffer rate of 500 Kbps is specified. The actual delay buffers allocated to each logical interface are 2 seconds of 500 Kbps. The 2-second value is based on the following calculation:

```
delay-buffer-rate < [8 x 64 Kbps]): 2 seconds of delay-buffer-rate
```

For more information about this calculation, see the Junos OS Class of Service Configuration Guide.

```
chassis {
  fpc 3 {
    pic 0 {
      q-pic-large-buffer;
    }
  }
}
interfaces {
  t1-3/0/1 {
    per-unit-scheduler;
  }
}
class-of-service {
  traffic-control-profiles {
    tc-profile3 {
      guaranteed-rate 750k;
      scheduler-map sched-map3;
      delay-buffer-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
    }
    tc-profile4 {
      guaranteed-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
      scheduler-map sched-map4;
    }
  }
  interface t1-3/0/1 {
    unit 0 {
      output-traffic-control-profile tc-profile3;
    }
    unit 1 {
      output-traffic-control-profile tc-profile4;
    }
  }
}
```


Configuring Link Services and CoS on Services PICs

To configure link services and CoS on an AS or Multiservices PIC, you must perform the following steps:

1. Enable the Layer 2 service package. You enable service packages per PIC, not per port. When you enable the Layer 2 service package, the entire PIC uses the configured package. To enable the Layer 2 service package, include the **service-package** statement at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services]** hierarchy level, and specify **layer-2**:

   ```
   [edit chassis fpc slot-number pic pic-number adaptive-services]
   service-package layer-2;
   ```

   For more information about AS or Multiservices PIC service packages, see Enabling Service Packages and "Layer 2 Service Package Capabilities and Interfaces" on page 3.

2. Configure a multilink PPP or FRF.16 bundle by combining constituent links into a virtual link, or bundle.

   #### Configuring an MLPPP Bundle

   To configure an MLPPP bundle, configure constituent links and bundle properties by including the following statements in the configuration:

   ```
   [edit interfaces interface-name unit logical-unit-number]
   encapsulation ppp;
   family mlppp {
     bundle lsq-fpc/pic/port.logical-unit-number;
   }
   [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;
   }
   ```

   For more information about these statements, see the Link and Multilink Properties.

   #### Configuring an MLFR FRF.16 Bundle

   To configure an MLFR FRF.16 bundle, configure constituent links and bundle properties by including the following statements in the configuration:

   ```
   [edit chassis fpc slot-number pic slot-number]
   mlfr-uni-nni-bundles number;
   [edit interfaces interface-name ]
   encapsulation multilink-frame-relay-uni-nni;
   unit logical-unit-number {
     family mlfr-uni-nni {
       bundle lsq-fpc/pic/port:channel;
   ```



```
    }
}
```

For more information about the **mlfr-uni-nni-bundles** statement, see the Junos OS System Basics Configuration Guide. MLFR FRF.16 uses channels as logical units.

For MLFR FRF.16, you must configure one end as data circuit-terminating equipment (DCE) by including the following statements at the **[edit interfaces lsq-*fpc/pic/port*:*channel*]** hierarchy level.

```
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
}
unit logical-unit-number {
    dlci dlci-identifier;
    family inet {
        address address;
    }
}
```

For more information about MLFR UNI NNI properties, see Link and Multilink Properties.

3. To configure CoS components for each multilink bundle, enable per-unit scheduling on the interface, configure a scheduler map, apply the scheduler to each queue, configure a fragmentation map, and apply the fragmentation map to each bundle. Include the following statements:

```
[edit interfaces]
lsq-fpc/pic/port {
    per-unit-scheduler; # Enables per-unit scheduling on the bundle
}
[edit class-of-service]
interfaces {
    lsq-fpc/pic/port { # Multilink PPP
    unit logical-unit-number {
        scheduler-map map-name; # Applies scheduler map to each queue
    }
}
lsq-fpc/pic/port:channel { # MLFR FRF.16
```



```
unit logical-unit-number {
   # Scheduler map provides scheduling information for
   # the queues within a single DLCI.
   scheduler-map map-name;
   shaping-rate percent percent;
}
forwarding-classes {
   queue queue-number class-name priority (high | low);
}
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 (percent percentage | rate | remainder) <exact>;
   }
}
fragmentation-maps {
   map-name {
      forwarding-class class-name {
         fragment-threshold bytes;
         no-fragmentation;
      }
   }
}
```

Associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI by including the following statements at the **[edit class-of-service]** hierarchy level:

```
interfaces {
   lsq-fpc/pic/port {
      unit logical-unit-number { # Multilink PPP
      fragmentation-map map-name;
   }
}
lsq-fpc/pic/port:channel { # MLFR FRF.16
unit logical-unit-number {
   fragmentation-map map-name;
}
```

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP

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;
```

NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ 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;
}
```

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, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP" on page 40.

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 | percent percentage | 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 Configuration Guide.

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 {
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
```

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. If you choose to set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For more information about MCML, see "Configuring Multiclass MLPPP on LSQ Interfaces" on page 25. For more information about fragmentation maps, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

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.

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 chassis]
```

```
fpc 1 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
t1-0/0/0 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1; # This adds t1-0/0/0 to the specified bundle.
    }
  }
}
t1-0/0/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0 {
  unit 1 { # This is the virtual link that concatenates multiple T1s.
  encapsulation multilink-ppp;
  drop-timeout 1000;
  fragment-threshold 128;
  link-layer-overhead 0.5;
  minimum-links 2;
  mrru 4500;
  short-sequence;
  family inet {
    address 10.2.3.4/24;
  }
}
[edit interfaces]
lsq-1/3/0 {
  per-unit-scheduler;
}
[edit class-of-service]
interfaces {
  lsq-1/3/0 { # multilink PPP constituent link
  unit 0 {
    scheduler-map sched-map1;
  }
}
t1-0/0/0 { # multilink PPP constituent link
unit 0 {
  scheduler-map sched-map1;
}
t1-0/0/1 { # multilink PPP constituent link
unit 0 {
  scheduler-map sched-map1;
}
```

```
forwarding-classes {
   queue 0 be;
   queue 1 ef;
   queue 2 af;
   queue 3 nc;
}
scheduler-maps {
   sched-map1 {
      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 30;
      buffer-size percent 30;
      priority low;
   }
   be-scheduler {
      transmit-rate percent 25;
      buffer-size percent 25;
      priority low;
   }
   ef-scheduler {
      transmit-rate percent 40;
      buffer-size percent 40;
      priority strict-high; # voice queue
   }
   nc-scheduler {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority high;
   }
}
fragmentation-maps {
   fragmap-1 {
      forwarding-class be {
         fragment-threshold 180;
      }
      forwarding-class ef {
         fragment-threshold 100;
      }
   }
}
[edit interfaces]
lsq-1/3/0 {
   unit 0 {
      fragmentation-map fragmap-1;
   }
}
```


Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16

To configure an *N*xT1 bundle using FRF.16, you aggregate *N* different T1 links into a bundle. The *N*xT1 bundle carries a potentially large number of Frame Relay PVCs, identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency.

To aggregate T1 links into an FRF.16 bundle, include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *slot-number*]** hierarchy level and include the **bundle** statement at the **[edit interfaces t1-*fpc/pic/port* unit *logical-unit-number* family mlfr-uni-nni]** hierarchy level:

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

[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlfr-uni-nni]
bundle lsq-fpc/pic/port:channel;
```

NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq- *fpc/pic/port*:*channel*]** hierarchy level:

```
[edit interfaces lsq- fpc/pic/port:channel]
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 number;
  t392 number;
  yellow-differential-delay milliseconds;
}
unit logical-unit-number {
  dlci dlci-identifier;
  family inet {
    address address;
```

```
    }
  }
```

The link services IQ channel represents the FRF.16 bundle. Four queues are associated with each DLCI. A scheduler removes packets from the queues according to a scheduling policy. On the link services IQ interface, you typically designate one queue to have strict priority. The remaining queues are serviced in proportion to weights you configure.

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 Configuration Guide.

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

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

For FRF.16, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16" on page 46.

For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bandwidth to be used by queue 0. For M Series and T Series routers, the default schedulers' transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent. These default schedulers send all user traffic to queue 0 and all network-control traffic to queue 3, and therefore are well suited to the behavior of FRF.16. If desired, you can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behavior, and apply it to the constituent links.

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.**

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

```
[edit class-of-service]
interfaces {
  lsq-fpc/pic/port:channel {
    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 | percent percentage | remainder) <exact>;
    }
  }
```

To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

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

For FRF.16 traffic, only multilink encapsulated (fragmented and sequenced) queues are supported. This is the default queuing behavior for all forwarding classes. FRF.16 does not allow for nonencapsulated traffic because the protocol requires that all packets carry the fragmentation header. If a large packet is split into multiple fragments, the fragments must have consecutive sequential numbers. Therefore, you cannot include the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level for FRF.16 traffic. For FRF.16, if you want to carry voice or any other latency-sensitive traffic, you should not use slow links. At T1 speeds and above, the serialization delay is small enough so that you do not need to use explicit LFI.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.16 header. The FRF.16 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*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** 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*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** 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.

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. Because each packet has an FRF.16 header, the sequence number field is used to put the packet back into sequence number order.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16

Configure an *N*xT1 bundle using FRF.16 with multiple CoS scheduler maps:

```
[edit chassis fpc 1 pic 3]
adaptive-services {
  service-package layer-2;
}
mlfr-uni-nni-bundles 2; # Creates channelized LSQ interfaces/FRF.16 bundles.
[edit interfaces]
t1-0/0/0 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle lsq-1/3/0:1;
    }
  }
}
t1-0/0/1 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle lsq-1/3/0:1;
    }
  }
}
lsq-1/3/0:1 { # Bundle link consisting of t1-0/0/0 and t1-0/0/1
per-unit-scheduler;
encapsulation multilink-frame-relay-uni-nni;
dce; # One end needs to be configured as DCE.
mlfr-uni-nni-bundle-options {
  drop-timeout 180;
  fragment-threshold 64;
  hello-timer 180;
  minimum-links 2;
  mrru 3000;
  link-layer-overhead 0.5;
}
```



```
unit 0 {
   dlci 26; # Each logical unit maps a single DLCI.
   family inet {
      address 10.2.3.4/24;
   }
}
unit 1 {
   dlci 42;
   family inet {
      address 10.20.30.40/24;
   }
}
unit 2 {
   dlci 69;
   family inet {
      address 10.20.30.40/24;
   }
}
[edit class-of-service]
scheduler-maps {
   sched-map-lsq0 {
      forwarding-class af scheduler af-scheduler-lsq0;
      forwarding-class be scheduler be-scheduler-lsq0;
      forwarding-class ef scheduler ef-scheduler-lsq0;
      forwarding-class nc scheduler nc-scheduler-lsq0;
   }
   sched-map-lsq1 {
      forwarding-class af scheduler af-scheduler-lsq1;
      forwarding-class be scheduler be-scheduler-lsq1;
      forwarding-class ef scheduler ef-scheduler-lsq1;
      forwarding-class nc scheduler nc-scheduler-lsq1;
   }
}
schedulers {
   af-scheduler-lsq0 {
      transmit-rate percent 60;
      buffer-size percent 60;
      priority low;
   }
   be-scheduler-lsq0 {
      transmit-rate percent 30;
      buffer-size percent 30;
      priority low;
   }
   ef-scheduler-lsq0 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority strict-high;
   }
   nc-scheduler-lsq0 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority high;
   }
   af-scheduler-lsq1 {
      transmit-rate percent 50;
```

```
      buffer-size percent 50;
      priority low;
    }
    be-scheduler-lsq1 {
      transmit-rate percent 30;
      buffer-size percent 30;
      priority low;
    }
    ef-scheduler-lsq1 {
      transmit-rate percent 15;
      buffer-size percent 15;
      priority strict-high;
    }
    nc-scheduler-lsq1 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority high;
    }
  }
  interfaces {
    lsq-1/3/0:1 { # MLFR FRF.16
    unit 0 {
      scheduler-map sched-map-lsq0;
    }
    unit 1 {
      scheduler-map sched-map-lsq1;
    }
  }
```

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI

When you configure a single fractional T1 interface, it is called a logical interface, because it can represent, for example, a routing adjacency.

The logical link services IQ interface represents the MLPPP bundle. Four queues are associated with the logical interface. 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.

To configure a single fractional T1 interface using MLPPP and LFI, you associate one DS0 (fractional T1) interface with a link services IQ interface. To associate a fractional T1 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-*fpc/pic/port*:*channel* unit *logical-unit-number* family mlppp]** hierarchy level:

[edit interfaces ds-*fpc/pic/port*:*channel* unit *logical-unit-number* family mlppp]
bundle lsq-*fpc/pic/port*.*logical-unit-number*;

NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ 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;
}
```

For MLPPP, assign a single scheduler map to the link services IQ (**lsq**) interface 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 (**lsq**) interface and to each constituent link and to each constituent link, as shown in "Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI" on page 51.

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.

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

```
[edit class-of-service]
interfaces {
  ds-fpc/pic/port.channel {
    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 | percent percentage | remainder) <exact>;
  }
```

```
}
```

For link services IQ interfaces, a strict-high-priority queue might starve all the other 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 receives infinite credits and does round-robin with high-priority queues, as described in the Junos OS Class of Service Configuration Guide.

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:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

When a packet is removed from a multilink-encapsulated queue, it is fragmented. 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.

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.

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 the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an MLPPP header, the software assumes the packet is a fragment of a larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI

Configure a single fractional T1 logical interface:

```
[edit interfaces]
lsq-0/2/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-ppp;
    link-layer-overhead 0.5;
    family inet {
      address 10.40.1.1/30;
    }
  }
}
ct3-1/0/0 {
  partition 1 interface-type ct1;
}
ct1-1/0/0:1 {
  partition 1 timeslots 1-2 interface-type ds;
}
ds-1/0/0:1:1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-0/2/0.0;
    }
  }
}
[edit class-of-service]
interfaces {
  ds-1/0/0:1:1 { # multilink PPP constituent link
  unit 0 {
    scheduler-map sched-map1;
```

```
  }
}
forwarding-classes {
  queue 0 be;
  queue 1 ef;
  queue 2 af;
  queue 3 nc;
}
scheduler-maps {
  sched-map1 {
    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 20;
    buffer-size percent 20;
    priority low;
  }
  be-scheduler {
    transmit-rate percent 20;
    buffer-size percent 20;
    priority low;
  }
  ef-scheduler {
    transmit-rate percent 50;
    buffer-size percent 50;
    priority strict-high; # voice queue
  }
  nc-scheduler {
    transmit-rate percent 10;
    buffer-size percent 10;
    priority high;
  }
}
fragmentation-maps {
  fragmap-1 {
    forwarding-class be {
      fragment-threshold 180;
    }
    forwarding-class ef {
      fragment-threshold 100;
    }
  }
}
[edit interfaces]
lsq-0/2/0 {
  unit 0 {
    fragmentation-map fragmap-1;
  }
}
```


Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12

To configure a single fractional T1 interface using FRF.16, you associate a DS0 interface with a link services IQ (**lsq**) interface. When you configure a single fractional T1, the fractional T1 carries a potentially large number of Frame Relay PVCs identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency. To associate the DS0 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-*fpc/pic/port*:*channel* unit *logical-unit-number* family mlfr-end-to-end]** hierarchy level:

```
[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlfr-end-to-end]
bundle lsq-fpc/pic/port.logical-unit-number;
```

NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ 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-frame-relay-end-to-end;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
  address address;
}
```

The logical link services IQ interface represents the FRF.12 bundle. Four queues are associated with each logical interface. 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 FRF.12, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. For M Series and T Series routers, the default schedulers, 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 FRF.12, you should configure schedulers with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign them to the link services IQ interface (**lsq**) and to each constituent link, as shown in "Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12" on page 56.

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.

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

```
[edit class-of-service]
interfaces {
  ds-fpc/pic/port.channel {
    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 | percent percentage | 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 Configuration Guide.

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:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see “Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 22.

When a packet is removed from a multilink-encapsulated queue, it is fragmented. 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.

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.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.12 header. The FRF.12 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 the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain Frame Relay header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an FRF.12 header, the software assumes the packet is a fragment of a larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain Frame Relay header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

A whole packet from a nonencapsulated queue can be placed between fragments of a multilink-encapsulated queue. However, fragments from one multilink-encapsulated queue cannot be interleaved with fragments from another multilink-encapsulated queue. This is the intent of the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*. If fragments from two different queues were interleaved, the header fields might not have enough information to separate the fragments.

Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12

FRF.12 with Fragmentation and Without LFI

This example shows a 128 KB DS0 interface. There is one traffic stream on **ge-0/0/0**, which is classified into queue 0 (**be**). Packets are fragmented in the link services IQ (**lsq-**) interface according to the threshold configured in the fragmentation map.

```
[edit chassis]
fpc 0 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.1.1.1/24 {
        arp 20.1.1.2 mac 00.90.1b.12.34.56;
      }
    }
  }
}
ce1-0/2/0 {
  partition 1 timeslots 1-2 interface-type ds;
}
ds-0/2/0:1 {
  no-keepalives;
  dce;
  encapsulation frame-relay;
  unit 0 {
    dlci 100;
    family mlfr-end-to-end {
      bundle lsq-0/3/0.0;
    }
  }
}
lsq-0/3/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 10.200.0.78/30;
    }
  }
}
fxp0 {
  unit 0 {
    family inet {
      address 172.16.1.162/24;
```

```
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
        }
    }
}
[edit class-of-service]
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    lsq-0/3/0 {
        unit 0 {
            fragmentation-map map1;
        }
    }
}
fragmentation-maps {
    map1 {
        forwarding-class {
            be {
                fragment-threshold 160;
            }
        }
    }
}
```

FRF.12 with Fragmentation and LFI

This example shows a 512 KB DS0 bundle and four traffic streams on **ge-0/0/0** that are classified into four queues. The fragment size is 160 for queue 0, queue 1, and queue 2. The voice stream on queue 3 has LFI configured.

```
[edit chassis]
fpc 0 {
    pic 3 {
        adaptive-services {
            service-package layer-2;
        }
    }
}
[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 20.1.1.1/24 {
                arp 20.1.1.2 mac 00.90.1b.12.34.56;
            }
```

```
    }
}
ce1-0/2/0 {
    partition 1 timeslots 1-8 interface-type ds;
}
ds-0/2/0:1 {
    no-keepalives;
    dce;
    encapsulation frame-relay;
    unit 0 {
        dlci 100;
        family mlfr-end-to-end {
            bundle lsq-0/3/0.0;
        }
    }
}
lsq-0/3/0 {
    per-unit-scheduler;
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
            address 10.200.0.78/30;
        }
    }
}
[edit class-of-service]
classifiers {
    inet-precedence ge-interface-classifier {
        forwarding-class be {
            loss-priority low code-points 000;
        }
        forwarding-class ef {
            loss-priority low code-points 010;
        }
        forwarding-class af {
            loss-priority low code-points 100;
        }
        forwarding-class nc {
            loss-priority low code-points 110;
        }
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    lsq-0/3/0 {
        unit 0 {
            scheduler-map sched2;
            fragmentation-map map2;
        }
    }
    ds-0/2/0:1 {
```



```
    scheduler-map link-map2;
  }
  ge-0/0/0 {
    unit 0 {
      classifiers {
        inet-precedence ge-interface-classifier;
      }
    }
  }
}
scheduler-maps {
  sched2 {
    forwarding-class be scheduler economy;
    forwarding-class ef scheduler business;
    forwarding-class af scheduler stream;
    forwarding-class nc scheduler voice;
  }
  link-map2 {
    forwarding-class be scheduler link-economy;
    forwarding-class ef scheduler link-business;
    forwarding-class af scheduler link-stream;
    forwarding-class nc scheduler link-voice;
  }
}
fragmentation-maps {
  map2 {
    forwarding-class {
      be {
        fragment-threshold 160;
      }
      ef {
        fragment-threshold 160;
      }
      af {
        fragment-threshold 160;
      }
      nc {
        no-fragmentation;
      }
    }
  }
  schedulers {
    economy {
      transmit-rate percent 26;
      buffer-size percent 26;
    }
    business {
      transmit-rate percent 26;
      buffer-size percent 26;
    }
    stream {
      transmit-rate percent 35;
      buffer-size percent 35;
    }
    voice {
      transmit-rate percent 13;
```

```
                buffer-size percent 13;
            }
            link-economy {
                transmit-rate percent 26;
                buffer-size percent 26;
            }
            link-business {
                transmit-rate percent 26;
                buffer-size percent 26;
            }
            link-stream {
                transmit-rate percent 35;
                buffer-size percent 35;
            }
            link-voice {
                transmit-rate percent 13;
                buffer-size percent 13;
            }
        }
    }
}
```

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15

This example configures an ***N***xT1 bundle using FRF.15 on a link services IQ interface. FRF.15 is similar to FRF.12, as described in "Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12" on page 53. The difference is that FRF.15 supports multiple physical links in a bundle, whereas FRF.12 supports only one physical link per bundle. For the Junos OS implementation of FRF.15, you can configure one DLCI per physical link.

NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. This example refers to T1 interfaces, but the configuration for E1 interfaces is similar.

```
[edit interfaces]
lsq-1/3/0 {
    per-unit-scheduler;
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
    }
}
unit 1 {
    encapsulation multilink-frame-relay-end-to-end;
}
# First physical link
t1-1/1/0:1 {
    encapsulation frame-relay;
    unit 0 {
        dlci 69;
        family mlfr-end-to-end {
            bundle lsq-1/3/0.0;
        }
```



```
      }
    }
    # Second physical link
    t1-1/1/0:2 {
      encapsulation frame-relay;
      unit 0 {
        dlci 13;
        family mlfr-end-to-end {
          bundle lsq-1/3/0.0;
        }
      }
    }
```

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

This example bundles a single T3 interface on a link services IQ interface with MLPPP encapsulation. Binding a single T3 interface to a multilink bundle allows you to configure compressed RTP (CRTP) on the T3 interface.

This scenario applies to MLPPP bundles only. The Junos OS does not currently support CRTP over Frame Relay. For more information, see Configuring Services Interfaces for Voice Services.

There is no need to configure LFI at DS3 speeds, because the packet serialization delay is negligible.

```
[edit interfaces]
t3-0/0/0 {
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0.1 {
  encapsulation multilink-ppp;
}
compression {
  rtp {
    # cRTP parameters go here
    #
    port minimum 2000 maximum 64009;
  }
}
```

This configuration uses a default fragmentation map, which results in all forwarding classes (queues) being sent out with a multilink header.

To eliminate multilink headers, you can configure a fragmentation map in which all queues have the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, and attach the fragmentation map to the **lsq-1/3/0.1** interface, as shown here:

```
[edit class-of-service]
fragmentation-maps {
```

```
    fragmap {
      forwarding-class {
        be {
          no-fragmentation;
        }
        af {
          no-fragmentation;
        }
        ef {
          no-fragmentation;
        }
        nc {
          no-fragmentation;
        }
      }
    }
  }
  interfaces {
    lsq-1/3/0.1 {
      fragmentation-map fragmap;
    }
  }
```

Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12

This example configures a clear-channel T3 or OC3 interface with multiple logical interfaces (DLCIs) on the link. In this scenario, each DLCI represents a customer. DLCIs are shaped at the egress PIC to a particular speed (*N*xDS0). This allows you to configure LFI using FRF.12 End-to-End Protocol on Frame Relay DLCIs.

To do this, first configure logical interfaces (DLCIs) on the physical interface. Then bundle the DLCIs, so that there is only one DLCI per bundle.

The physical interface must be capable of per-DLCI scheduling, which allows you to attach shaping rates to each DLCI. For more information, see the Junos® OS Network Interfaces.

To prevent fragment drops at the egress PIC, you must assign a shaping rate to the link services IQ logical interfaces and to the egress DLCIs. Shaping rates on DLCIs specify how much bandwidth is available for each DLCI. The shaping rate on link services IQ interfaces should match the shaping rate assigned to the DLCI that is associated with the bundle.

Egress interfaces also must have a scheduler map attached. The queue that carries voice should be strict-high-priority, while all other queues should be low-priority. This makes LFI possible.

This example shows voice traffic in the **ef** queue. The voice traffic is interleaved with bulk data. Alternatively, you can use multiclass MLPPP to carry multiple classes of traffic in different multilink classes, as described in "Configuring Multiclass MLPPP on LSQ Interfaces" on page 25.

```
[edit interfaces]
```



```
t3-0/0/0 {
  per-unit-scheduler;
  encapsulation frame-relay;
  unit 0 {
    dlci 69;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.0;
    }
  }
  unit 1 {
    dlci 42;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0 {
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
  }
  fragment-threshold 320; # Multilink packets must be fragmented
}
unit 1 {
  encapsulation multilink-frame-relay-end-to-end;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
  sched { # Scheduling parameters that apply to bundles on AS or Multiservices PICs.
  ...
}
pic-sched {
  # Scheduling parameters for egress DLCIs.
  # The voice queue should be strict-high priority.
  # All other queues should be low priority.
  ...
}
fragmentation-maps {
  fragmap {
    forwarding-class {
      ef {
        no-fragmentation;
      }
      # Voice is carried in the ef queue.
      # It is interleaved with bulk data.
    }
  }
}
interfaces {
  t3-0/0/0 {
    unit 0 {
      shaping-rate 512k;
      scheduler-map pic-sched;
    }
    unit 1 {
      shaping-rate 128k;
```



```
        scheduler-map pic-sched;
      }
    }
    lsq-1/3/0 { # Assign fragmentation and scheduling to LSQ interfaces.
    unit 0 {
      shaping-rate 512k;
      scheduler-map sched;
      fragmentation-map fragmap;
    }
    unit 1 {
      shaping-rate 128k;
      scheduler-map sched;
      fragmentation-map fragmap;
    }
  }
```

For more information about how FRF.12 works with links services IQ interfaces, see "Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12" on page 53.

Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP

This example configures an ATM2 IQ interface with MLPPP bundled with link services IQ interfaces. This allows you to configure LFI on ATM virtual circuits.

For this type of configuration, the ATM2 IQ interface must have LLC encapsulation.

The following ATM PICs are supported in this scenario:

- 2-port OC-3/STM1 ATM2 IQ
- 4-port DS3 ATM2 IQ

Virtual circuit multiplexed PPP over AAL5 is not supported. Frame Relay is not supported. Bundling of multiple ATM VCs into a single logical interface is not supported.

Unlike DS3 and OC3 interfaces, there is no need to create a separate scheduler map for the ATM PIC. For ATM, you define CoS components at the **[edit interfaces at-*fpc/pic/port* atm-options]** hierarchy level, as described in the Junos® OS Network Interfaces.

NOTE: **Do not configure RED profiles on ATM logical interfaces that are bundled. Drops do not occur at the ATM interface.**

In this example, two ATM VCs are configured and bundled into two link services IQ bundles. A fragmentation map is used to interleave voice traffic with other multilink traffic. Because MLPPP is used, each link services IQ bundle can be configured for CRTP.

```
[edit interfaces]
at-1/2/0 {
  atm-options {
    vpi 0;
    pic-type atm2;
  }
```

```
    unit 0 {
        vci 0.69;
        encapsulation atm-mlppp-llc;
        family mlppp {
            bundle lsq-1/3/0.10;
        }
    }
    unit 1 {
        vci 0.42;
        encapsulation atm-mlppp-llc;
        family mlppp {
            bundle lsq-1/3/0.11;
        }
    }
}
lsq-1/3/0 {
    unit 10 {
        encapsulation multilink-ppp;
    }
    # Large packets must be fragmented.
    # You can specify fragmentation for each forwarding class.
    fragment-threshold 320;
    compression {
        rtp {
            port minimum 2000 maximum 64009;
        }
    }
}
unit 11 {
    encapsulation multilink-ppp;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
    sched { # Scheduling parameters that apply to LSQ bundles on AS or Multiservices PICs.
    ...
}
fragmentation-maps {
    fragmap {
        forwarding-class {
            ef {
                no-fragmentation;
            }
        }
    }
}
interfaces { # Assign fragmentation and scheduling parameters to LSQ interfaces.
lsq-1/3/0 {
    unit 0 {
        shaping-rate 512k;
        scheduler-map sched;
        fragmentation-map fragmap;
    }
    unit 1 {
        shaping-rate 128k;
        scheduler-map sched;
```

```
        fragmentation-map fragmap;
    }
}
```


CHAPTER 3

Configuration Statements

cisco-interoperability

Syntax cisco-interoperability send-lip-remove-link-for-link-reject;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description FRF.16 interoperability settings.

Options **send-lip-remove-link-for-link-reject**—Send Link Integrity Protocol remove link when an add-link rejection message is received.

Usage Guidelines See "Configuring SONET APS Interoperability with Cisco Systems FRF.16" on page 9.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

forwarding-class

Syntax

```
forwarding-class class-name {
  (fragment-threshold bytes | no-fragmentation);
  multilink-class number;
}
```

Hierarchy Level [edit class-of-service fragmentation-maps]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define a forwarding class name and associated fragmentation properties within a fragmentation map.

The **fragment-threshold** and **no-fragmentation** statements are mutually exclusive.

Default If you do not include this statement, the traffic in forwarding class ***class-name*** is fragmented.

Options ***class-name***—Name of the forwarding class.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

fragment-threshold

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, set the fragmentation threshold for an individual forwarding class.

Default If you do not include this statement, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** 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 maximum transmission unit (MTU) of all the links in the bundle.

Options ***bytes***—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

Range: 128 through 16,320 bytes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 22

fragmentation-map

Syntax fragmentation-map *map-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI.

Default If you do not include this statement, traffic in all forwarding classes is fragmented.

Options ***map-name***—Name of the fragmentation map.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

fragmentation-maps

Syntax

```
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      (fragment-threshold bytes | no-fragmentation);
      multilink-class number;
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define fragmentation properties for individual forwarding classes.

Default If you do not include this statement, traffic in all forwarding classes is fragmented.

Options *map-name*—Name of the fragmentation map.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 22.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

hot-standby

Syntax hot-standby;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options],
[edit interfaces rlsq*number:number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description For one-to-one AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds. For FRF.15 (MLFR) and FRF.16 (MFR) configuration, specify the switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

link-layer-overhead

Syntax link-layer-overhead *percent*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead. Link-layer overhead accounts for the bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information. Overhead resulting from link-layer encapsulation and framing is computed automatically.

Options *percent*—Percentage of total bundle bandwidth to be set aside for link-layer overhead.

Range: 0 through 50 percent

Default: 0 percent

Usage Guidelines See "Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 19.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

lsq-failure-options

Syntax

```
lsq-failure-options {
  no-termination-request;
  trigger-link-failure interface-name;
}
```

Hierarchy Level [edit interfaces lsq-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define the failure recovery option settings.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring the Association between LSQ and SONET Interfaces" on page 8.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

multilink-class

Syntax multilink-class *number*;

Hierarchy Level [edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, map a forwarding class into a multiclass MLPPP (MCML).

The **multilink-class** statement and **no-fragmentation** statements are mutually exclusive.

Options *number*—The multilink class assigned to this forwarding class.

Range: 0 through 7

Default: None

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 22
- Configuring Multiclass MLPPP on LSQ Interfaces on page 25
- Configuring Fragmentation by Forwarding Class
- Junos Services Interfaces Configuration Release 12.3
- multilink-max-classes on page 306

multilink-max-classes

Syntax multilink-max-classes *number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, configure the number of multilink classes to be negotiated when a link joins the bundle.

Options *number*—The number of multilink classes to be negotiated when a link joins the bundle.

Range: 1 through 8

Default: None

Usage Guidelines See "Configuring Multiclass MLPPP on LSQ Interfaces" on page 25.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-fragmentation

Syntax no-fragmentation;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, set traffic on a particular forwarding class to be interleaved, rather than fragmented. This statement specifies that no extra fragmentation header is prepended to the packets received on this queue and that static-link load balancing is used to ensure in-order packet delivery.

Static-link load balancing is done based on packet payload. For IP version 4 (IPv4) and IP version 6 (IPv6) traffic, the link is chosen based on a hash computed from the source address, destination address, and protocol. If the IP payload is Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) traffic, the hash also includes source port and destination port. For MPLS traffic, the hash includes all MPLS labels and fields in the payload, whether the MPLS payload is IPv4 or IPv6.

Default If you do not include this statement, the traffic in forwarding class ***class-name*** is fragmented.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 22

no-per-unit-scheduler

Syntax no-per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 11.4.

Description To enable traffic control profiles to be applied at FRF.16 bundle (physical) interface level, disable the per-unit scheduler, which is enabled by default. This statement and the **shared-scheduler** statement are mutually exclusive.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Oversubscribing Interface Bandwidth

no-termination-request

Syntax no-termination-request;

Hierarchy Level [edit interfaces *interface-name* ppp-options],
[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* ppp-options]** hierarchy level added in Junos OS Release 8.3.

Description Inhibit PPP termination-request messages to the remote host if the primary circuit fails.

Usage Guidelines See "Configuring the Association between LSQ and SONET Interfaces" on page 8.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For channelized OC3 IQ, channelized OC12 IQ, channelized STM1 IQ, channelized T3 IQ, channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), link services (ls-) on J Series routers, Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-Gigabit Ethernet interfaces only, enable association of scheduler map names with logical interfaces.

NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.

NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the **per-unit-scheduler** statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the **per-unit-scheduler** statement, the entire PIC supports 1024 – 2 * number of ports (1024 minus two times the number of ports), because each port is allocated two default schedulers.

When including the **per-unit-scheduler** statement for interfaces on the IQ2 and IQ2-E PIC, you must also include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- vlan-tagging

preserve-interface

Syntax preserve-interface;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in Junos OS Release 7.6.

Description Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:

- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC

Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.

Usage Guidelines See "Configuring Link State Replication for Redundant Link PICs" on page 13.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces

primary

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the primary Link Services IQ PIC interface.

Options *interface-name*—The identifier for the Link Services IQ PIC interface, which must be of the form **lsq-*fpc/pic/port***.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

redundancy-options

Syntax

```
redundancy-options {
  (hot-standby | warm-standby);
  primary lsq-fpc/pic/port;
  secondary lsq-fpc/pic/port;
}
```

Hierarchy Level [edit interfaces rlsq*number*]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the primary and secondary (backup) Link Services IQ PIC interfaces.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

secondary

Syntax secondary *interface-name*;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the secondary (backup) Link Services IQ PIC interface.

Options ***interface-name***—The identifier for the Link Services IQ PIC interface, which must be of the form **lsq-*fpc/pic/port***.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

trigger-link-failure

Syntax trigger-link-failure *interface-name*;

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the Link Services IQ PIC fails.

Options ***interface-name***—Name of SONET interface.

Usage Guidelines See “Configuring the Association between LSQ and SONET Interfaces” on page 8.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

warm-standby

Syntax warm-standby;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 8.0.

Description For AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery involves one backup PIC supporting multiple working PICs. Recovery time is not guaranteed.

Usage Guidelines See “Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces” on page 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

PART 3

Administration

CHAPTER 4

Link Services Operational Mode Commands

show services link-services cpu-usage

Syntax

show services link-services cpu-usage
<brief | detail>
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 8.4.

Description

(M Series and T Series routers only) Display information about Link Services IQ (LSQ) CPU usage.

Options

none—Display standard information about CPU usage for all LSQ interfaces.

brief | detail—(Optional) Display the specified level of output.

interface ***interface-name***—(Optional) Display information about the specified LSQ interface.

Required Privilege Level

view

List of Sample Output

show services link-services cpu-usage brief (AS PIC) on page 84
show services link-services cpu-usage brief (MultiServices PIC) on page 84
show services link-services cpu-usage detail (AS PIC) on page 84
show services link-services cpu-usage detail (MultiServices PIC) on page 85

Output Fields

Table 3 on page 82 lists the output fields for the **show services link-services cpu-usage** command. Output fields are listed in the approximate order in which they appear.

Table 3: show services link-services cpu-usage Output Fields

Field Name	Field Description	Level of Output
Role	CPU functional category.	**brief**
1 Second Average	Percentage of usage during 1-second duration.	All levels
5 Second Average	Percentage of usage during 5-second duration.	All levels
QoS	Quality of service (QoS) CPU, which takes care of queuing and scheduling of incoming IP packets on a per-bundle basis. It schedules packets with higher QoS values first.	All levels
Sequencer	Assigns sequence numbers to outgoing MLPPP fragments and interleaves link fragmentation and interleaving (LFI) traffic.	All levels
Load Balancer	Distributes load across different fragmenter CPUs.	All levels
Fragmenter	Main LSQ CPU; fragments IP packets into MLPPP fragments and also reassembles MLPPP fragments into IP packets.	All levels
Total	Sum of all CPU functions.	**brief**

Table 3: show services link-services cpu-usage Output Fields *(continued)*

Field Name	Field Description	Level of Output
Idle	Counts idle cycles when the CPU does not have any work.	**detail**
Timer	Takes care of periodic events driven by a timer, such as timeouts.	**detail**
System	System housekeeping thread.	**detail**
Input (QoS)	Acquires and queues incoming IP frames from hardware interfaces.	**detail**
Output (QoS)	Sends scheduled frames to the next processing CPU.	**detail**
Output Frags (QoS)	Sends outstanding frames to the fragmenter CPU.	**detail**
Bypass (QoS)	Sends outstanding frames for LFI.	**detail**
Free frame (QoS)	Frees dropped frames.	**detail**
CPU*number*	Identifier number of specific CPU.	**detail**
Drop (Fragmenter)	Drops frames that have been marked by the QoS CPU.	**detail**
Frag (Fragmenter)	Fragments IP frames into MLPPP fragments.	**detail**
Reass (Fragmenter)	Reassembles MLPPP fragments into IP frames.	**detail**
Freeback (Fragmenter)	Handles freeback of credits from other CPUs (MultiServices PICs only).	**detail**
Input LFI (Sequencer)	Receives LFI traffic from QoS CPU and transmits it with strict priority over MLPPP.	**detail**
Input Frag (Sequencer)	Receives MLPPP fragments from fragmenter CPUs, assigns sequence numbers, and appends MLPPP headers.	**detail**
Output Frag (Sequencer)	Load-balances and transmits fragments across links.	**detail**
Retry (Sequencer)	Retries transmission if hardware was busy in the previous attempt.	**detail**
Input Alloc (Load Balancer)	Acquires frames from hardware interfaces and validates them.	**detail**
Input (Load Balancer)	Performs error and sanity checks and check frames for PortMapping.	**detail**
Output (Load Balancer)	Sends frame to next processing CPU.	**detail**

Table 3: show services link-services cpu-usage Output Fields *(continued)*

Field Name	Field Description	Level of Output
Freeback (Load Balancer)	Handles freeback of credits from other CPUs.	**detail**

Sample Output

show services link-services cpu-usage brief (AS PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
Role                1 Second Average       5 Second Average
QOS                             1.0%                   1.0%
Sequencer                      0.1%                   0.1%
Fragmenter                      0.1%                   0.1%
Total                           0.1%                   0.1%
```

show services link-services cpu-usage brief (MultiServices PIC)

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 brief
  Role              1 Second Average       5 Second Average

  QoS                           0.1%                   0.1%
  Fragmenter                    0.1%                   0.1%
  Load Balancer                 0.0%                   0.0%
  Total                         0.1%                   0.1%
```

show services link-services

```
user@host> show services link-services cpu-usage interface lsq-0/0/0 detail

  QoS                Idle   Timer   System  Input Output Output Bypass  Free
```

cpu-usage detail (AS PIC)

```
                                                        Frags           frame
CPU0            99.1%     0.9%     0.0%    0.0%    0.0%    0.0%    0.0%    0.0%
CPU1            99.8%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%    0.0%
1 sec ave       99.5%     0.5%     0.0%    0.0%    0.0%    0.0%    0.0%    0.0%
5 sec ave       99.5%     0.5%     0.0%    0.0%    0.0%    0.0%    0.0%    0.0%

Fragmenter       Idle    Timer   System    Drop    Frag   Reass    Free
                                                                   back

CPU0            96.6%     0.1%     0.0%    0.0%    0.0%    3.3%    0.0%
CPU1            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU2            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU3            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU4            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU5            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU6            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU7            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
CPU8            99.9%     0.1%     0.0%    0.0%    0.0%    0.0%    0.0%
1 sec ave       99.5%     0.1%     0.0%    0.0%    0.0%    0.4%    0.0%
5 sec ave       99.5%     0.1%     0.0%    0.0%    0.0%    0.4%    0.0%

Sequencer        Idle    System   Input   Input Output   Retry
                                    LFI    Frag   Frag

CPU0            99.9%      0.1%    0.0%    0.0%    0.0%    0.0%
CPU1           100.0%      0.0%    0.0%    0.0%    0.0%    0.0%
1 sec ave       99.9%      0.1%    0.0%    0.0%    0.0%    0.0%
5 sec ave       99.9%      0.1%    0.0%    0.0%    0.0%    0.0%
```

show services link-services

user@host> **show services link-services cpu-usage interface lsq-0/0/0 detail**

```
QoS              Idle    Timer   System   Input Output Output Bypass   Free
                                                        Frags          frame
```

cpu-usage detail (MultiServices PIC)

```
CPU0            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
CPU1            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
CPU2            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
CPU3            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
CPU4            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
1 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%
5 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%   0.0%

Fragmenter       Idle   Timer  System   Drop   Frag  Reass   Free
                                                             back

CPU0            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU1            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU2            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU3            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU4            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU5            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU6            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU7            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU8            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU9            99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU10           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU11           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU12           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU13           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU14           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU15           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU16           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
CPU17           99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
1 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%
5 sec ave       99.9%    0.1%    0.0%   0.0%   0.0%   0.0%   0.0%

Load-Balancer    Idle   System  Input  Input Output   Free
                                Alloc                 back

CPU0           100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
CPU1           100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
1 sec ave      100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
5 sec ave      100.0%     0.0%   0.0%   0.0%   0.0%   0.0%
```

CHAPTER 5

Link Services Interface Operational Mode Commands

show interfaces (Link Services IQ)

Syntax

show interfaces lsq-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<l2-statistics>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
l2-statistics option introduced with Junos OS Release 12.1.

Description

(J Series, M Series, MX Series, and T Series routers only) Display status information about the specified link services intelligent queuing (IQ) interface.

Options

lsq-*fpc/pic/port*—Display standard status information about the specified link services IQ interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

l2-statistics—(Optional) Display Layer 2 queue statistics for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Additional Information

Link services IQ interfaces are similar to link services interfaces. The important difference is that link services IQ interfaces fully support Junos OS class-of-service (CoS) components.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (MLPPP on Link Services IQ) on page 103
show interfaces extensive (Multiclass MLPPP on Link Services IQ) on page 104
show interfaces extensive (MLPPP on Link Services IQ Bundle) on page 106
show interfaces extensive (MFR on Link Services IQ Bundle) on page 108
show interfaces (Multiclass MLPPP on Link Services IQ) on page 110

Output Fields

Table 4 on page 89 lists the output fields for the **show interfaces** (link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 4: show interfaces (Link Services IQ) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: **Multilink-Frame-Relay-UNI-NNI** (default), **LinkService**, **Frame-relay**, **Frame-relay-ccc**, or **Frame-relay-tcc**.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second timezone*** **(*hour*:*minute*:*second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Multilink Frame Relay UNI NNI bundle options	Multilink Frame Relay UNI NNI only) Configured information about Multilink Frame Relay bundle options.
	• **Device type**—**DCE** (data communication equipment) or **DTE** (data terminal equipment).
	• **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** to **4500** bytes. The default is **1524** bytes.
	• **Bandwidth**—Speed at which the interface is running.
	• **Fragmentation threshold**—Configured fragmentation threshold: **128** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation.
	• **Red differential delay limit**—Red differential delay limit among bundle links has been reached, indicating an action will occur.
	• **Yellow differential delay limit**—Yellow differential delay among bundle links has been reached, indicating a warning will occur.
	• **Red differential delay action**—Type of actions taken when the red differential delay exceeds the red limit: **Disable link transmit** or **Remove link from service**.
	• **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link layer overhead.
	• **Reassembly drop timer**—Drop timeout value to provide a recovery mechanism if individual links in the link services bundle drop one or more packets: **1** through **127** milliseconds. By default, the drop timeout parameter is **0** (disabled). A value under **5** ms is not recommended.
	• **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**.
	• **LIP Hello timer**—Link Interleaving Protocol hello timer: **1** through **180** seconds.
	  ◦ **Acknowledgement timer**—Maximum period to wait for an add link acknowledgement, hello acknowledgement, or remove link acknowledgement: **1** through **10** seconds.
	  ◦ **Acknowledgement retries**—Number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgement timer: **1** through **5.**
		detail extensive none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Multilink Frame Relay UNI NNI bundle options (continued)	• **Bundle class**—Bundle class ID.
	• **LMI type**—Multilink Frame Relay UNI NNI LMI type:**ANSI**, **Q.933 ANNEX A**, or **Consortium**.
	  ◦ **T391 LIV polling timer**—Multilink Frame Relay UNI NNI Full status polling counter: **1** through **255**, with a default value of **6**.
	  ◦ **T392 polling verification timer**—Multilink Frame Relay UNI NNI LMI error threshold. The number of errors required to bring down the link, within the event count specified by **N393**. The range is **1** through **10**, with a default value of **3**.
	  ◦ **N391 full status polling count**—Multilink Frame Relay UNI NNI Full status polling counter: **1** through **255**.
	  ◦ **N392 error threshold**—Multilink Frame Relay UNI NNI LMI error threshold: **1** through **10**.
	• **N393 monitored event count**—Multilink Frame Relay UNI NNI LMI monitored event count: **1** through **10**, with a default value of **4.**
	• **Consortium LMI Settings**
	  ◦ **n391dte**—DTE full status polling interval in seconds: **1** through **255**.
	  ◦ **n392dce**—DCE error threshold: **1** through **10**.
	  ◦ **n392dte**—DTE error threshold: **1** through **10**.
	  ◦ **n393dce**—DCE monitored event count: **1** through **10**. **n393dte**—DTE monitored event count: **1** through **10**.
	  ◦ **t391dte**—DTE polling verification timer (in seconds): **5** through **30**.
	  ◦ **t392dce**—DCE polling verification timer (in seconds): **5** through **30**.
		detail extensive none

LMI	Local Managment Interface packet statistics: • **Input**—Number of packets arriving on the interface (nn) and timestamp of the most recent packet arrival, in the format: **Input:** *nn* **(last seen** *hh:mm:ss ago***)** **Output**—Number of packets sent out on the interface (nn) and how much time has passed since the last packet was sent, in the format: **Output:** *nn* **(last seen** *hh:mm:ss ago***)**	**detail extensive** none
DTE Statistics	Statistics about information transferred from the data terminal equipment (DTE) to the data communications equipment (DCE). • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DCE from the DTE. • **Full enquiry responses received**—Number of full enquiry responses received by DCE from the DTE.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
DCE Statistics	Statistics about information transferred from the DCE to the DTE. • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none
Common Statistics	Statistics about messages snet between the DTE and the DCE. • **Unknown messages received**—Number of received packets that do not fall into any other category. • **Asynchronouts updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timed out**—Number of keepalive responses that time out when no Local Management Interface (LMI) packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the Packet Forwarding Engine (PFE). Input traffic refers to the fragments received by the ingress PFE, which get assembled into Layer 3 input packets. Output packets refer to the IP packets transmitted out of the ingress PFE to the LSQ, which get segmented into output fragments.	**detail extensive**

DLCI*nn*	Data-link connection identifier (DLCI) number of the logicial interview. The following information is displayed.
	• **Flags**—Values are:
	  ◦ **Active**—Set when the link is active and the DTE and DCE are exchanging information.
	  ◦ **Down**—Set when the link is active, but no information is received from the DTE.
	  ◦ **DCE unconfigured**—Set when the corresponding DCLI in the DCE is not configured.
	  ◦ **Configured**—Set when the correspondening DLCCI is configured.
	  ◦ **DCE-Configured**—Displayed when the command is issued from the DTE.	

DLCI Statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	
Input rate	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified
Output rate	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**
Frame exceptions	Information about framing exceptions. Includes events recorded under **Exception Events** for each logical interface. • **Oversized frames**—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • **Errored input frames**—Number of input frame errors. • **Input on disabled link/bundle**—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • **Output for disabled link/bundle**—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • **Queuing drops**—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**extensive**
Buffering exceptions	Information about buffering exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Assembly exceptions	(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under **Exception Events** for each logical interface. An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**
Hardware errors (sticky)	(Multilink Frame Relay end-to-end only) Information about hardware errors: • **Data memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • **Control memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Queue counters	Queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive** none
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
Encapsulation	Encapsulation being used: PPP or Multilink PPP.	All levels
Bandwidth	Speed at which the interface is running.	All levels
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or **Multilink class status** (Multiclass MLPPP)	Information about bundle status: • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • **Received sequence number**—Sequence number for received packets. • **Transmitted sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully, but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Statistics	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	
	• **Bundle**—Information for each active bundle link.
	  ◦ **Fragments: Input and Output**—Total number and rate of fragments received and transmitted.
	  ◦ **Packets: Input and Output**—Total number and rate of packets received and transmitted.
	  ◦ **Multilink class**—(Multiclass MLPPP only) Information about multiclass links used in the multilink operation.
	• **Link**—Information about links used in the multilink operation.
	  ◦ **Link name**—Interface name of the link services IQ channel and state information (physical link **up** or **down**).
	  ◦ **Input and Output**—Total number and rate of fragments and packets received and transmitted.
		detail extensive

NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	**detail extensive** none
Addresses, Flags	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
MLPPP Bundle Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
SNMP-Traps	SNMP trap notifications are enabled.	All levels
Encapsulation	Encapsulation being used: PPP, Multilink PPP, or Multilink-FR.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Bandwidth	Speed at which the interface is running.	All levels
Bundle links Iinformation	Information about the bundled links. • **Active bundle links**—Number of active links. • **Removed bundle links**—Information about links used in the multilink operation. • **Disabled bundle links**—Number of disabled links.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Inner PPP Protocol field compression**—Inner PPP protocol compression is enabled or disabled. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP)	Information about bundle status: • **Received sequence number**—Sequence number for received packets. • **Transmit sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	**detail extensive** none

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Statistics	Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	
	The bundle, multilink, and network statistics are reported by the Packet Forwarding Engine (PFE). The Multi Link Detail statistics like fragments, non-fragments and LFI are reported by the PIC.
	
	However, the PFE reports an extra overhead of 2 bytes in the output when compared with the Multilink Detail Statistics. This is due to the service-cookie in the PFE which does the link demux for the ML header.
	
	The difference in the bytes received and transmitted from Network and Multilink interfaces and Multilink statistics for each member link is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows.
	
	• Input side - Total overhead = 6 bytes.
	  ◦ ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.
	  ◦ PPP: 2 bytes of protocol field.
	• Output side - Total overhead = 11 bytes.
	  ◦ ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.
	  ◦ PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.
	  ◦ 2 bytes of Service Cookie.
	
	• **Bundle**—Information for each active bundle link.
	  ◦ **Multilink: Input and Ouput**—Total number and rate of multilink frames, bytes, and bits per second received and transmitted. It is a module connecting LSQ PIC and its member link. Multilink Input displays L2 fragments received from the member link to the LSQ PIC. Multilink Output displays the L2 fragments transmitted from LSQ PIC to the member links.
	  ◦ **Network: Input and Output**—Total number of network frames, bytes, and bits per second received and transmitted. It refers to the packets transmitted from an ingress interface to the PFE and then to the LSQ PIC. Network Input displays the L3 packets received from the LSQ PIC to the PFE. Network Output displays the L3 packets transmitted from PFE to LSQ PIC.
	• **Link**—Information about links used in the multilink operation.
	  ◦ **Link name**—The interface name of the link services IQ channel and state information (physical link **up** or **down**) and up time.
	  ◦ **Input and Output**—Total number and rate of frames, bytes, and bits per second received and transmitted.
		extensive

Table 4: show interfaces (Link Services IQ) Output Fields *(continued)*

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>Multilink detail statistics</td><td>Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).

The difference in the bytes received and transmitted from the bundle is divided between the ML and the PPP headers. For example the header counter for a long sequence configuration would be as follows:

• Input side - Total overhead = 6 bytes.
  ◦ ML: 4 bytes of ML header = 1 byte of Flag + 3 bytes of long sequence number.
  ◦ PPP: 2 bytes of protocol field.
• Output side - Total overhead = 9 bytes.
  ◦ ML: 4 bytes of ML Header = 1 byte of Flag + 3 bytes of Long sequence number.
  ◦ PPP: 5 bytes = 4 bytes of header + 1 byte of Idle flag.

• Bundle—Information for the bundle link.
  ◦ Fragments: Input and Output—Total number and rate of multilink fragments received and transmitted.
  ◦ Non-fragments: Input and Output—Total number and rate of nonfragmented multilink frames received and transmitted.
  ◦ LFI: Input and Output—Total number and rate of link fragmented and interleaved frames and bytes.</td><td>extensive</td></tr>
<tr><td>Protocol</td><td>Protocol family configured on the logical interface.</td><td>detail extensive none</td></tr>
<tr><td>MTU</td><td>MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked Adjusted.</td><td>detail extensive none</td></tr>
<tr><td>Generation</td><td>Unique number for use by Juniper Networks technical support only.</td><td>detail extensive</td></tr>
<tr><td>Route Table</td><td>Routing table in which this address exists. For example, Route table:0 refers to inet.0.</td><td>detail extensive</td></tr>
<tr><td>Addresses, Flags</td><td>Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.</td><td>detail extensive none</td></tr>
<tr><td>Destination</td><td>IP address of the remote side of the connection.</td><td>detail extensive none</td></tr>
<tr><td>Local</td><td>IP address of the logical interface.</td><td>detail extensive none</td></tr>
<tr><td>Broadcast</td><td>Broadcast address on the logical interface.</td><td>detail extensive none</td></tr>
<tr><td>Generation</td><td>Unique number for use by Juniper Networks technical support.</td><td>detail extensive</td></tr>
</table>

Sample Output

show interfaces extensive (MLPPP on Link Services IQ)

```
user@host> show interfaces lsq-0/2/0 extensive
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:05:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              8872424               229080 bps
   Output bytes  :              9856960               234448 bps
   Input  packets:                38202                  117 pps
   Output packets:                39453                  117 pps
  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
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 be                           0                    0                    0

    1 ef                           0                    0                    0

    2 af                           0                    0                    0

    3 nc                           0                    0                    0

  Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 256kbps
    Bundle options:
      MRRU                           1504
      Drop timer period              2000
      Sequence number format         long (24 bits)
      Fragmentation threshold        0
      Links needed to sustain bundle 1
      Multilink classes              0
      Link layer overhead            4.0 %
    Bundle status:
      Remote MRRU                    1500
      Received sequence number       0x0
      Transmit sequence number       0x0
      Packet drops                   0 (0 bytes)
      Fragment drops                 9 (1401 bytes)
      MRRU exceeded                  0
```



```
    Fragment timeout              0
    Missing sequence number       0
    Out-of-order sequence number  4
    Out-of-range sequence number  0
    Packet data buffer overflow   0
    Fragment data buffer overflow 0
  Statistics         Frames        fps         Bytes          bps
  Bundle:
    Multilink:
      Input :         79827        239       9593009       232288
      Output:         77533        234       9811743       238056
    Network:
      Input :         38202        117       8872424       229080
      Output:         39453        117       9856960       234448
  Link:
    ds-1/0/2:1:1.0 <-- up
      Input :          1114         87        180183       113608
      Output:          1577        118        199215       119064
    ds-1/0/2:1:2.0 <-- down
      Input :          1941        152        187948       118680
      Output:          1574        116        199494       118992
  Protocol inet, MTU: 1500 [Adjusted]
    Flags: User-MTU, MTU-Protocol-Adjusted
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.74.11/24, Local: 10.74.11.10
  Protocol iso, MTU: 1500 [Adjusted]
    Flags: User-MTU, MTU-Protocol-Adjusted
  Protocol mpls, MTU: 1488 [Adjusted], Maximum labels: 3
    Flags: User-MTU, MTU-Protocol-Adjusted
```

show interfaces extensive (Multiclass

```
user@host> show interfaces extensive lsq-0/2/0
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
```

MLPPP on Link Services IQ)

```
Link-level type: LinkService, MTU: 1504
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :              3474024              223704 bps
 Output bytes  :              4193992              233888 bps
 Input  packets:                15809                 116 pps
 Output packets:                16788                 117 pps
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
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 be                            0                    0                    0

  1 ef                            0                    0                    0

  2 af                            0                    0                    0

  3 nc                            0                    0                    0

Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU                           1504
    Drop timer period              2000
    Sequence number format         long (24 bits)
    Fragmentation threshold        0
    Links needed to sustain bundle 1
    Multilink classes              2
    Link layer overhead            4.0 %
  Multilink class 0 status:
    Received sequence number       0x4c38
    Transmit sequence number       0x4890
    Packet drops                   0 (0 bytes)
    Fragment drops                 2551 (397084 bytes)
    MRRU exceeded                  0
    Fragment timeout               52
    Missing sequence number        0
    Out-of-order sequence number   953
    Out-of-range sequence number   0
    Packet data buffer overflow    0
    Fragment data buffer overflow  0
  Multilink class 1 status:
```

```
    Received sequence number        0xffffff
    Transmit sequence number        0x3710
    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:
    Fragments:
      Input :          33719        239       4041763       231632
      Output:          32371        234       4096545       237488
    Packets:
      Input :          15809        116       3474024       223704
      Output:          16788        117       4193992       233888
  Multilink class 0:
    Fragments:
      Input :          19331          0             0            0
      Output:              0          0             0            0
    Packets:
      Input :           2064          0             0            0
      Output:           1864          0             0            0
  Multilink class 1:
    Fragments:
      Input :              0          0             0            0
      Output:          14096          0             0            0
    Packets:
      Input :          14096          0             0            0
      Output:              0          0             0            0
  Link:
    ds-1/0/2:1:1.0, Enabled, Physical link is Up
      Input :          20972        151       2030595       118080
      Output:          16184        116       2048468       118488
    ds-1/0/2:1:2.0, Enabled, Physical link is Up
      Input :          12747         88       2011168       113552
      Output:          16187        118       2048077       119000
  Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
    Flags: User-MTU, MTU-Protocol-Adjusted
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
      Generation: 18
```

show interfaces extensive (MLPPP on

```
user@host> show interfaces lsq-7/1/0.0 extensive
Logical interface lsq-7/1/0.0 (Index 88) (SNMP ifIndex 114) (Generation 188)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-FR
```

Link Services IQ Bundle)

```
  Last flapped: Never
  Bandwidth: 256kbps
  Bundle links information:
    Active bundle links       2
    Removed bundle links      0
    Disabled bundle links     0
  Bundle options:
    MRRU                                 1504
    Drop timer period                    1500
    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 %
  Bundle status:
    Received sequence number        0xb74
    Transmit sequence number        0xb74
    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 :         315381          0      42757818            0
      Output:         315381          0      43388580            0
    Network:
      Input :         315381          0      40952064            0
      Output:         315381          0      40952064            0
  Link:
    ds-6/0/0:1:1.0
      Up time: Up since boot
      Input :          63794          0      25146728            0
      Output:          63778          0      25273164            0
    ds-6/0/0:1:2.0
      Up time: Up since boot
      Input :         251587          0      17611090            0
      Output:         251603          0      18115416            0
  Multilink detail statistics:
  Bundle:
    Fragments:
      Input :              0          0             0            0
      Output:              0          0             0            0
    Non-fragments:
      Input :         293748          0      19387368            0
      Output:         293748          0      20562360            0
    LFI:
      Input :          21633          0      22152192            0
      Output:          21633          0      22325256            0
  Protocol inet, MTU: 1500, Generation: 204, Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast:
Unspecified, Generation: 214
```

show interfaces extensive (MFR on Link Services IQ Bundle)

```
user@host> show interfaces lsq-1/0/0:0 extensive
Physical interface: lsq-1/0/0:0, Enabled, Physical link is Up
  Interface index: 179, SNMP ifIndex: 746, Generation: 182
  Link-level type: Multilink-FR-UNI-NNI, MTU: 1508
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Last flapped   : 2010-11-15 01:11:00 PST (00:31:58 ago)
  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                         1536kbps
    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                          Consortium
      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
  Consortium LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
  LMI statistics:
    Input : 188 (last seen 00:00:01 ago)
    Output: 189 (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                   : 157
    Full enquiries received              : 31
    Enquiry responses sent               : 158
    Full enquiry responses sent          : 31
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 0
  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       1
    Removed bundle links      0
    Disabled bundle links     0
Multilink Frame Relay UNI NNI active bundle links statistics:
                   Frames        fps         Bytes          bps
  t1-7/0/0:1:3.0
    Up time: 00:31:24
    Input :            0          0             0            0
    Output:            0          0             0            0
    Current differential delay       0.0 ms
    Recent high differential delay   0.0 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       2       0        0     189       0       0
    Xmt:      2       1       0      189       0       0       0

Logical interface lsq-1/0/0:2.0 (Index 77) (SNMP ifIndex 751) (Generation 142)

  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-FR-UNI-NNI
  Last flapped: 2010-11-15 01:11:40 PST (00:31:18 ago)
  Bundle status:
    Received sequence number        0xfff
    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 :            0          0             0            0
      Output:            0          0             0            0
    Network:
      Input :            0          0             0            0
      Output:            0          0             0            0
  Link:
    t1-7/0/0:1:3.0
      Up time: 00:31:24
      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: 153, Route table: 0
        Flags: Sendbcast-pkt-to-re
        Addresses, Flags: Is-Preferred Is-Primary
          Destination: 10.0.1.8/30, Local: 10.0.1.9, Broadcast: Unspecified,
Generation: 154
      DLCI 12
        Flags: Active
        Total down time: 00:00:32 sec, Last down: 00:31:50 ago
        Traffic statistics:
         Input  bytes  :                    0
         Output bytes  :                    0
         Input  packets:                    0
         Output packets:                    0
    DLCI statistics:
      Active DLCI  :1  Inactive DLCI  :0
```

show interfaces (Multiclass MLPPP on Link Services IQ)

user@host> **show interfaces extensive lsq-0/2/0**

```
Physical interface: lsq-0/2/0, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 25, Generation: 23
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2005-06-02 08:54:36 PDT (00:02:25 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              3474024              223704 bps
   Output bytes  :              4193992              233888 bps
   Input  packets:                15809                 116 pps
   Output packets:                16788                 117 pps
  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
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 be                            0                    0                    0

    1 ef                            0                    0                    0

    2 af                            0                    0                    0

    3 nc                            0                    0                    0
```

```
Logical interface lsq-0/2/0.0 (Index 66) (SNMP ifIndex 26) (Generation 5)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
  Bandwidth: 256kbps
  Bundle options:
    MRRU                          1504
    Drop timer period             2000
    Sequence number format        long (24 bits)
    Fragmentation threshold       0
    Links needed to sustain bundle  1
    Multilink classes             2
    Link layer overhead           4.0 %
  Multilink class 0 status:
    Received sequence number      0x4c38
    Transmit sequence number      0x4890
    Packet drops                  0 (0 bytes)
    Fragment drops                2551 (397084 bytes)
    MRRU exceeded                 0
    Fragment timeout              52
    Missing sequence number       0
    Out-of-order sequence number  953
    Out-of-range sequence number  0
    Packet data buffer overflow   0
    Fragment data buffer overflow 0
  Multilink class 1 status:
    Received sequence number      0xffffff
    Transmit sequence number      0x3710
    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:
    Fragments:
      Input :         33719       239      4041763       231632
      Output:         32371       234      4096545       237488
    Packets:
      Input :         15809       116      3474024       223704
      Output:         16788       117      4193992       233888
  Multilink class 0:
    Fragments:
      Input :         19331         0            0            0
      Output:             0         0            0            0
    Packets:
      Input :          2064         0            0            0
      Output:          1864         0            0            0
  Multilink class 1:
    Fragments:
      Input :             0         0            0            0
      Output:         14096         0            0            0
    Packets:
      Input :         14096         0            0            0
      Output:             0         0            0            0
  Link:
    ds-1/0/2:1:1.0, Enabled, Physical link is Up
      Input :         20972       151      2030595       118080
```

```
    Output:         16184        116       2048468       118488
  ds-1/0/2:1:2.0, Enabled, Physical link is Up
    Input :         12747         88       2011168       113552
    Output:         16187        118       2048077       119000
Protocol inet, MTU: 1500 [Adjusted], Generation: 14, Route table: 0
  Flags: User-MTU, MTU-Protocol-Adjusted
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.0.1.0/30, Local: 10.0.1.2, Broadcast: Unspecified,
    Generation: 18
```


show interfaces (Redundant Link Services IQ)

Syntax

show interfaces rlsq*number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<queue>
<routing>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 7.6.

Description

(M Series and T Series routers only) Display status information about the specified redundant link services intelligent queuing (IQ) configuration.

Options

rlsq*number*—Redundant link services IQ interface name. The logical interface number range of values is **0** through **127**.

none—Display standard status information about the specified redundant link services IQ configuration.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

queue—(Optional) Display queue information about network interfaces.

routing—(Optional) Display routing information about network interfaces.

snmp-index ***snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Redundant Link Services IQ) on page 125
show interfaces brief (Redundant Link Services IQ) on page 125
show interfaces detail (Redundant Link Services IQ) on page 125
show interfaces extensive (Redundant Link Services IQ) on page 127

Output Fields

Table 5 on page 113 lists the output fields for the **show interfaces** (redundant link services IQ) command. Output fields are listed in the approximate order in which they appear.

Table 5: show interfaces (Redundant Link Services IQ) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: **Multilink-Frame-Relay-UNI-NNI** (default), **LinkService**, **Frame-relay**, **Frame-relay-ccc**, or **Frame-relay-tcc**.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second timezone*** **(*hour*:*minute*:*second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input rate	(Redundant LSQ) Rate of bits and packets received on the interface.	None specified
Output rate	(Redundant LSQ) Rate of bits and packets transmitted on the interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.	**detail extensive**

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Frame exceptions	Information about framing exceptions. Includes events recorded under **Exception Events** for each logical interface. • **Oversized frames**—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb (kilobits). • **Errored input frames**—Number of input frame errors. • **Input on disabled link/bundle**—Number of frames received on disabled links. These frames can result either from an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it. • **Output for disabled link/bundle**—Number of frames sent for a disabled or unavailable link. These frames can result either from an inconsistent configuration, or from a bundle being brought up or down while traffic is flowing through it. • **Queuing drops**—Total number of packets dropped before traffic enters the link services IQ interface. Indicates that the interface is becoming oversubscribed.	**extensive**
Buffering exceptions	Information about buffering exceptions. Includes events recorded under **Exception Events** for each logical interface: • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Assembly exceptions	(Multilink Frame Relay end-to-end only) Information about assembly exceptions. Includes events recorded under **Exception Events** for each logical interface. An assembly exception does not necessarily indicate an operational problem with the physical link services IQ interface itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the interface records one or more exception events. The physical interface can drop multilink-encapsulated fragments itself as a result. Any multilink packets or fragments dropped by the interface itself result in packet or fragment drop counts on individual logical interfaces. If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, excess differential delay or traffic losses on individual interfaces can be the root cause. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the physical link services IQ interface itself, but can occur when one or more individual links drop traffic. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the link services IQ interface unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible. • **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.	**extensive**
Hardware errors (sticky)	(Multilink Frame Relay end-to-end only) Information about hardware errors: • **Data memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support. • **Control memory error**—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks technical support.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Queue counters	Queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive** none
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
Encapsulation	Encapsulation being used: PPP or Multilink PPP.	All levels
Bandwidth	Speed at which the interface is running.	All levels
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP) or **Multilink class status** (MC-MLPPP)	Information about bundle status: • **Remote MRRU**—MRRU value received from remote peer. If negotiation has not been initiated, the default value is displayed. • **Received sequence number**—Sequence number for received packets. • **Transmitted sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—Received a frame with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Statistics	Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for each active bundle link.
	  ◦ **Fragments: Input and Output**—Total number and rate of fragments received and transmitted.
	  ◦ **Packets: Input and Output**—Total number and rate of packets received and transmitted.
	  ◦ **Multilink class**—(MC-MLPPP only) Information about multiclass links used in the multilink operation.
	• **Link**—Information about links used in the multilink operation.
	  ◦ **Link name**—Interface name of the link services IQ channel and state information (physical link **up** or **down**).
	  ◦ **Input and Output**—Total number and rate of fragments and packets received and transmitted.
		detail extensive

NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	**detail extensive** none
Addresses, Flags	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
MLPPP Bundle Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	All levels
SNMP-Traps	SNMP trap notifications are enabled.	All levels
Encapsulation	Encapsulation being used: PPP, Multilink PPP or Multilink-FR.	All levels
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second timezone*** **(*hour*:*minute*:*second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Bandwidth	Speed at which the interface is running.	All levels
Bundle links Iinformation	Information about the bundled links. • **Active bundle links**—Number of active links. • **Removed bundle links**—Information about links used in the multilink operation. • **Disabled bundle links**—Number of disabled links.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle options	(Multilink Frame Relay end-to-end interfaces only) • **MRRU**—Configured size of the maximum received reconstructed unit (MRRU): **1500** through **4500** bytes. The default is **1504** bytes. • **Drop timer period**—Drop timeout value to provide a recovery mechanism if individual links in link services bundle drop one or more packets: **0** though **2000** milliseconds. Values under 5 ms are not recommended. The default setting is **0**, which disables the timer. • **Inner PPP Protocol field compression**—Inner PPP protocol compression is enabled or disabled. • **Sequence number format**—Short sequence number header format (MLPPP only). • **Fragmentation threshold**—Configured fragmentation threshold: **64** through **16,320** bytes, in integer multiples of **64** bytes. The default setting is **0**, which disables fragmentation. • **Links needed to sustain bundle**—Minimum number of links to sustain the bundle: **1** through **8**. • **Multilink classes**—Number of multilink classes negotiated. • **Link layer overhead**—Percentage of bundle bandwidth to be set aside for link-layer overhead.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bundle status (MLPPP)	Information about bundle status: • **Received sequence number**—Sequence number for received packets. • **Transmit sequence number**—Sequence number for transmitted packets. • **Packet drops**—Number and byte count of output packets that were dropped, rather than being encapsulated and sent out of the router as fragments. The packet drop counter is incremented if there is a temporary shortage of packet memory on the AS PIC, which causes packet fragmentation to fail. • **Fragment drops**—Number and byte count of input fragments that were dropped, rather than being reassembled and handled by the router as packets. This counter also includes fragments that have been received successfully but had to be dropped because not all fragments that constituted a packet had been received. The fragment drop counter is incremented when a fragment received on constituent links is dropped. Drop fragments can be triggered by sequence ordering errors, duplicate fragments, timed-out fragments, and bad multilink headers. • **MRRU exceeded**—Number of reassembled packets exceeding the MRRU. This counter is not implemented in this release. • **Fragment timeout**—The drop timer expired while a fragment sequence number was outstanding. Occurs only if the drop timer is enabled. This timeout can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. • **Missing sequence number**—A gap was detected in the sequence numbers of fragments on a bundle. • **Out-of-order sequence number**—Two frames with out-of-order sequence numbers occurred within a single link. This event indicates that an individual link within a bundle reordered traffic, making the multilink interface unable to correctly process the resulting stream. • **Out-of-range sequence number**—A frame was received with an out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events can occur when the far end of a bundle is taken down or brought up. • **Packet data buffer overflow**—Packet buffer memory is full. This overflow can occur when the aggregate data rate exceeds the physical link services IQ interface capacity. • **Fragment data buffer overflow**—Fragment buffer memory is full. This overflow can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data rate exceeds the physical link services IQ capacity.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Statistics	Information about frames, bytes, and bits per second received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for each active bundle link.
	  • **Multilink: Input and Ouput**—Total number and rate of multilink frames, bytes, and bits per second received and transmitted.
	  • **Network: Input and Output**—Total number of multilink frames, bytes, and bits per second received and transmitted.
	• **Link**—Information about links used in the multilink operation.
	  • **Link name** is the interface name of the link services IQ channel and state information (physical link **up** or **down**) and up time.
	  • **Input and Output**—Total number and rate of frames, bytes, and bits per second received and transmitted.
		extensive

Multilink detail statistics	Frames, bytes, and bits per second received and sent by the bundle. All references to traffic direction (input or output) are defined with respect to the router. Each field has columns that indicate the number of frames received and transmitted, frames per second (fps), the number of bytes received and transmitted, and bits per second (bps).
	• **Bundle**—Information for the bundle link.
	  • **Fragments: Input and Output**—Total number and rate of multilink fragments received and transmitted.
	  • **Non-fragments: Input and Output**—Total number and rate of nonfragmented multilink frames received and transmitted.
	  • **LFI: Input and Output**—Total number and rate of link fragmented and interleaved frames and bytes.
		extensive

Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface. If the MTU value is negotiated down to meet the MRRU requirement on the remote side, this value is marked **Adjusted**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Addresses, Flags	Information about the addresses configured on the logical interface. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address on the logical interface.	**detail extensive** none

Table 5: show interfaces (Redundant Link Services IQ) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support.	**detail extensive**

Sample Output

show interfaces (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Bandwidth: 0
    Statistics         Frames        fps          Bytes          bps
    Bundle:
      Fragments:
        Input :             3          0            255            0
        Output:             3          0            264            0
      Packets:
        Input :             3          0            252            0
        Output:             0          0              0            0
    Link:
      t1-1/3/0:1.0
        Input :             3          0            255            0
        Output:             0          0              0            0
      t1-1/3/0:2.0
        Input :             0          0              0            0
        Output:             3          0            264            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2.2.2.0/30, Local: 2.2.2.1
```

show interfaces brief (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0 brief
Physical interface: rlsq0, Enabled, Physical link is Up
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Logical interface rlsq0.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    inet  2.2.2.1/30
```

show interfaces detail (Redundant Link Services IQ)

```
user@host> show interfaces rlsq0 detail
Physical interface: rlsq0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 27, Generation: 144
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 252                    0 bps
   Output bytes  :                 276                    0 bps
```



```
 Input  packets:                   3                    0 pps
 Output packets:                   3                    0 pps
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 be                            0                    0                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont                  0                    0                    0


Logical interface rlsq0.0 (Index 72) (SNMP ifIndex 88) (Generation 31)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Bandwidth: 0
  Bundle options:
    MRRU                           1504
    Remote MRRU                    N/A
    Drop timer period              2000
    Sequence number format         long (24 bits)
    Fragmentation threshold        0
    Links needed to sustain bundle  1
    Multilink classes              0
    Link layer overhead            4.0 %
  Bundle status:
    Received sequence number       0xffffff
    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:
    Fragments:
      Input :             3          0           255            0
      Output:             3          0           264            0
    Packets:
```



```
        Input :             3          0           252            0
        Output:             0          0             0            0
    Link:
      t1-1/3/0:1.0
        Input :             3          0           255            0
        Output:             0          0             0            0
      t1-1/3/0:2.0
        Input :             0          0             0            0
        Output:             3          0           264            0
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
    Protocol inet, MTU: 1500, Generation: 43, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2.2.2.0/30, Local: 2.2.2.1, Broadcast: Unspecified,
        Generation: 45
```

show interfaces extensive (Redundant Link Services IQ)

The output for the **show interfaces rlsq extensive** command is identical to that for the **show interfaces rlsq detail** command. For sample output, see show interfaces detail (Redundant Link Services IQ) on page 125.

PART 4

Index

Index

W