



Ethernet OAM Feature Guide for MX Series Routers



Published: 2013-08-12

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Ethernet OAM Feature Guide for MX Series Routers
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	ix
	Documentation and Release Notes	ix
	Supported Platforms	ix
	Using the Examples in This Manual	ix
	Merging a Full Example	x
	Merging a Snippet	x
	Documentation Conventions	xi
	Documentation Feedback	xiii
	Requesting Technical Support	xiii
	Self-Help Online Tools and Resources	xiii
	Opening a Case with JTAC	xiv
Part 1	Overview	
Chapter 1	IEEE 802.1ag OAM Connectivity-Fault Management	3
	Ethernet Operations, Administration, and Maintenance	3
	Ethernet OAM Connectivity Fault Management	4
Chapter 2	ITU-T Y.1731 Ethernet Frame Delay Measurements	7
	Ethernet Frame Delay Measurements	7
Chapter 3	IEEE 802.1ah OAM Link-Fault Management	11
	Ethernet OAM Link Fault Management	11
Chapter 4	Ethernet Ring Protection	13
	Ethernet Ring Protection	13
	Ethernet Ring Protection Using Ring Instances for Load Balancing	15
Part 2	Configuration	
Chapter 5	ITU-T Y.1731 Ethernet Frame Delay Measurements Tasks	19
	Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements	19
	Triggering an Ethernet Frame Delay Measurements Session	20
	Viewing Ethernet Frame Delay Measurements Statistics	21
Chapter 6	IEEE 802.1ag OAM Connectivity-Fault Management Examples	25
	Ethernet Operations, Administration, and Maintenance	25
	Ethernet OAM Connectivity Fault Management	26
	Example: Configuring Ethernet CFM over VPLS	27
	Example: Configuring Ethernet CFM on Bridge Connections	34
	Example: Configuring Ethernet CFM on Physical Interfaces	38

Chapter 7	ITU-T Y.1731 Ethernet Frame Delay Measurements Examples	41
	Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces	41
	Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces	46
	Example: Configuring Ethernet Frame Delay Measurements with Untagged Interfaces	50
Chapter 8	IEEE 802.1ah OAM Link-Fault Management Examples	53
	Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge	53
	Example: Configuring Ethernet LFM for CCC	54
	Example: Configuring Ethernet LFM for Aggregated Ethernet	56
	Example: Configuring Ethernet LFM with Loopback Support	58
Chapter 9	Ethernet Ring Protection Examples	61
	Example: Configuring Ethernet Ring Protection for MX Series Routers	61
	Example Topology	61
	Router 1 (RPL Owner) Configuration	63
	Router 2 Configuration	64
	Router 3 Configuration	66
	Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers	68
	Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation . . .	85
	Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition	87
Part 3	Administration	
Chapter 10	Managing Ethernet Frame Delay Measurements	93
	Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements . .	93
	Triggering an Ethernet Frame Delay Measurements Session	94

List of Figures

Part 1	Overview	
Chapter 2	ITU-T Y.1731 Ethernet Frame Delay Measurements	7
	Figure 1: Ethernet OAM Overview	8
Part 2	Configuration	
Chapter 6	IEEE 802.1ag OAM Connectivity-Fault Management Examples	25
	Figure 2: Ethernet OAM with VPLS	28
	Figure 3: Ethernet CFM over a Bridge Network	34
	Figure 4: Ethernet CFM on Physical Interfaces	38
Chapter 8	IEEE 802.1ah OAM Link-Fault Management Examples	53
	Figure 5: Ethernet LFM Between Provider Edge and Customer Edge	53
	Figure 6: Ethernet LFM for CCC	55
	Figure 7: Ethernet LFM for Aggregated Ethernet	56
	Figure 8: Ethernet LFM with Loopback Support	58
Chapter 9	Ethernet Ring Protection Examples	61
	Figure 9: Ethernet Ring Protection Example Nodes	62
	Figure 10: ERP with Multiple Protection Instances Configured on Three MX Series Routers	70

List of Tables

	About the Documentation	ix
	Table 1: Notice Icons	xi
	Table 2: Text and Syntax Conventions	xi
Part 2	Configuration	
Chapter 5	ITU-T Y.1731 Ethernet Frame Delay Measurements Tasks	19
	Table 3: Monitor Ethernet Delay Command Parameters	21
	Table 4: Show Ethernet Delay Command Parameters	22
Chapter 9	Ethernet Ring Protection Examples	61
	Table 5: Components of the Network Topology	70
Part 3	Administration	
Chapter 10	Managing Ethernet Frame Delay Measurements	93
	Table 6: Monitor Ethernet Delay Command Parameters	95

About the Documentation

- Documentation and Release Notes on page ix
- Supported Platforms on page ix
- Using the Examples in This Manual on page ix
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiii

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:

- 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.xml;
    }
  }
}
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.xml; }
```

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*)

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> 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.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	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.	broadcast multicast <i>(string1 string2 string3)</i>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

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

- [IEEE 802.1ag OAM Connectivity-Fault Management on page 3](#)
- [ITU-T Y.1731 Ethernet Frame Delay Measurements on page 7](#)
- [IEEE 802.1ah OAM Link-Fault Management on page 11](#)
- [Ethernet Ring Protection on page 13](#)

CHAPTER 1

IEEE 802.1ag OAM Connectivity-Fault Management

- [Ethernet Operations, Administration, and Maintenance on page 3](#)
- [Ethernet OAM Connectivity Fault Management on page 4](#)

Ethernet Operations, Administration, and Maintenance

This topic provides an overview to help you effectively configure Ethernet Operations, Administration, and Maintenance (OAM) on a network of Juniper Networks MX Series 3D Universal Edge Routers. For more information about configuring OAM parameters on Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual local area network (VLAN) identifier for troubleshooting
- Work independently of the actual Ethernet transport and function over physical Ethernet ports, or a virtual service such as pseudowire, and so on.
- Isolate faults over a flat (or single operator) network architecture or a nested or hierarchical (or multi-provider) networks.

OAM can provide simple link-level information, provide performance statistics, or track end-to-end connectivity across the network. Simple link fault management (LFM) for Ethernet links is defined in IEEE 802.3ah.

Ethernet OAM functions are implemented as:

- Fault detection and notification (provided by continuity check messages)
- Path discovery (provided by the linktrace protocol)
- Fault isolation, verification, and recovery (isolation and verification are provided by a combination of protocols, while recovery is the function of protocols such as spanning tree)

The loopback protocol used in Ethernet OAM is modeled on the standard IP ping. After a fault is detected, the loopback protocol performs fault verification and isolation under the direction of a network operator. The loopback is performed using request and response message pairs. A unicast loopback message is generated by a MEP and a loopback reply is generated by the destination MIP or MEP. The target MAC address is learned by the continuity check protocol or linktrace protocol. The loopback message's packet is always forwarded to a unique port by the originating MEP, as determined by a MAC table lookup or the MEP interface MAC address. The target MIP or MEP generates a unicast loopback reply in response to the received loopback message. The loopback message follows the same path as a data packet, and intermediate bridges simply forward the packet to the destination MIP or MEP.

**Related
Documentation**

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet OAM Connectivity Fault Management on page 4](#)
- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Ethernet OAM Connectivity Fault Management

The most complete connectivity fault management (CFM) is defined in IEEE 802.1ag. This topic emphasizes the use of CFM in a Metro Ethernet environment.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol which discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.
- Fault isolation using the loopback protocol. Similar to IP ping, this protocol works with the continuity check protocol during troubleshooting.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible. Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outermost domains are assigned a higher level than the innermost domains. Customer end points have to highest maintenance domain level. In a CFM maintenance domain, each service instance is called a maintenance association. A maintenance association can be thought as a full mesh of maintenance endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding

to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

MEPs can be *up MEPs* or *down MEPs*. A link can connect a MEP at level 5 to a MEP at level 7. The interface at level 5 is an up MEP (because the other end of the link is at MEP level 7) and the interface at level 7 is a down MEP (because the other end of the link is at MEP level 5).

In a Metro Ethernet network, CFM is commonly used at two levels:

- By the service provider to check the connectivity among its provider edge (PE) routers
- By the customer to check the connectivity among its customer edge (CE) routers



NOTE: The configured customer CFM level must be greater than service provider CFM level.

In many Metro Ethernet networks, CFM is used to monitor connectivity over a VPLS and bridge network.

**Related
Documentation**

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Operations, Administration, and Maintenance on page 3](#)
- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

CHAPTER 2

ITU-T Y.1731 Ethernet Frame Delay Measurements

- [Ethernet Frame Delay Measurements on page 7](#)

Ethernet Frame Delay Measurements

Performance management depends on the accurate measurement of service agreement objective parameters, which can include bandwidth and reliability. In many cases, a service provider could be subject to penalties imposed by regulation, statute, or contract if network performance is not within the bounds established for the service. One key performance objective is delay, along with its close relative, delay variation (often called jitter). Some applications will function just as well with high delays across the network and high delay variations (such as bulk file transfer), while other applications (such as voice) can only function with low and stable delays. Many networks invoke protocols or features available at Layer 3 (the packet layer) or higher to measure network delays and jitter link-by-link. However, when the network consists of many Ethernet links, there is little available at Layer 2 (the frame layer) that allows routers to measure frame delay and jitter. This is where the ability to configure and monitor Ethernet frame delay is helpful.

On a Juniper Networks MX Series Ethernet Services Router equipped with the Distributed Port Concentrator (MX-DPC) only, you can perform Ethernet frame delay measurements (referred to as ETH-DM in Ethernet specifications). This feature allows you to configure on-demand Operation, Administration, and Maintenance (OAM) statements for the measurement of frame delay and frame delay variation (jitter). You can configure Ethernet frame delay measurement in either one-way or two-way (round-trip) mode to gather frame delay statistics and simultaneous statistics from multiple sessions. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor Service Level Agreements (SLAs).

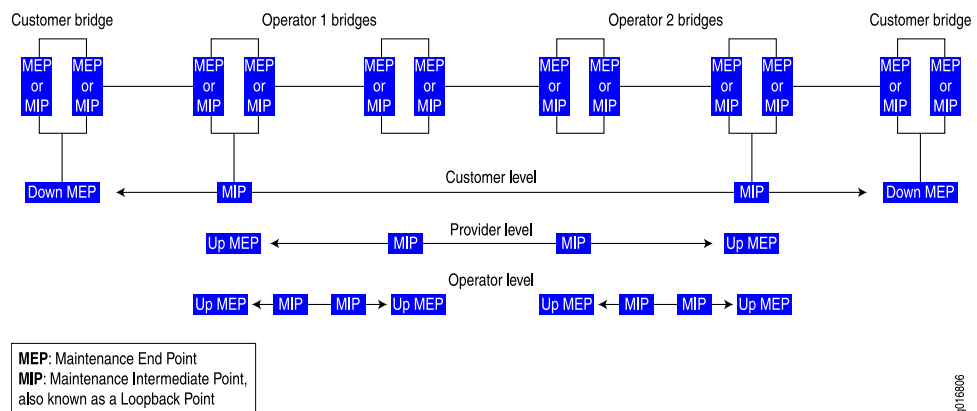
Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. Ethernet frame delay measurement supports hardware-based timestamping in the receive direction for delay measurements. It also provides runtime display of delay statistics when two-way delay measurement is triggered. Ethernet frame delay measurement records the last 100 samples collected per remote maintenance end point (MEP) or per connectivity fault management (CFM) session. You can retrieve the history at any time using simple commands. You can clear all Ethernet frame delay measurement statistics and PDU

counters. Ethernet frame delay measurement is fully compliant with the ITU-T Y.1731 (*OAM Functions and Mechanisms for Ethernet-based Networks*) specification.

Ethernet frame delay measurement uses the IEEE 802.1ag CFM infrastructure.

An overview of the architecture established for Ethernet OAM is shown in [Figure 1 on page 8](#). Generally, Ethernet frame delay measurements are made in a peer fashion from one MEP or CFM session to another. However, these measurements are not made to Maintenance Intermediate Points (MIPs).

Figure 1: Ethernet OAM Overview



There are two types of Ethernet frame delay measurements:

- One-way
- Two-way (round-trip)

For one-way Ethernet frame delay measurement, either MEP can send a request to begin a one-way delay measurement to its peer MEP. However, the statistics are collected only at the receiver MEP. This feature requires the clocks at the transmitting and receiving MEPs to be synchronized. If these clocks fall out of synchronization, only one-way delay variation and average delay variation values are computed correctly (and therefore valid). Use the `show` commands at the receiver MEP to display one-way delay statistics.

For two-way (round-trip) Ethernet frame delay measurement, either MEP can send a request to begin a two-way delay measurement to its peer MEP, which responds with timestamp information. Run-time statistics are collected and displayed at the initiator MEP. The clocks do not need to be synchronized at the transmitting and receiving MEPs. The Junos OS supports the optional timestamps in delay measurement reply (DMR) frames to increase the accuracy of delay calculations. The Junos OS also supports hardware-assisted timestamping for Ethernet frame delay protocol data units (PDUs) in the reception path.

Use the `show` commands at the initiator MEP to display two-way delay statistics, and at the receiver MEP to display one-way delay statistics.

The following are some limitations with regard to using Ethernet frame delay measurement:

- This feature is available only on MX Series routers.
- Ethernet frame delay measurements are available only when the distributed periodic packet management daemon (ppmd) is enabled.
- The statistics collected are lost after graceful Routing Engine switchover (GRES).
- You can monitor only one session to the same remote MEP or MAC address.
- Accuracy is compromised when the system changes (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.
- The use of Ethernet frame delay measurements on aggregated Ethernet and pseudowire interfaces is not supported.
- The use of hardware-assisted timestamping is not supported on all MX DPCs (Rev-B or higher is required).
- If you attempt to perform Ethernet frame delay measurements to a non-MX partner, the incoming Ethernet frame delay PDUs are discarded silently. Ethernet delay measurement commands and capabilities are not available on non-MX Series routers.

**Related
Documentation**

- *Ethernet OAM Feature Guide for MX Series Routers*
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
- [Triggering an ETH-DM Session on page 20](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

CHAPTER 3

IEEE 802.1ah OAM Link-Fault Management

- [Ethernet OAM Link Fault Management on page 11](#)

Ethernet OAM Link Fault Management

Link Fault Management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across a point-to-point Ethernet link either directly connected or through repeaters.

LFM provides the following functions:

- Failure detection on physical links in both directions, as well as unidirectional failures.
- Ability to put a port in link-loopback mode remotely for diagnostics.
- Report and receive link error events such as framing or symbol errors.

LFM runs at the physical or aggregated interface level. When configured on an aggregated interface, LFM is run individually on each member link. LFM is a link-layer protocol and does not need a Layer 3 (IPv4 or IPv6) address to operate. This allows for LFM to function on circuit cross-connect/transport cross-connect (CCC/TCC) encapsulated interfaces.

Related Documentation

- *Ethernet OAM Feature Guide for MX Series Routers*
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 53](#)
- [Example: Configuring Ethernet LFM for CCC on page 54](#)
- [Example: Configuring Ethernet LFM for Aggregated Ethernet on page 56](#)
- [Example: Configuring Ethernet LFM with Loopback Support on page 58](#)

CHAPTER 4

Ethernet Ring Protection

- [Ethernet Ring Protection on page 13](#)
- [Ethernet Ring Protection Using Ring Instances for Load Balancing on page 15](#)

Ethernet Ring Protection

Link failure is often an unavoidable part of networking. However, there are methods of improving the reliability of a router or bridge network even when link failures occur. For example, SONET/SDH seal-healing rings are frequently used to add a level of robustness to router networks. This ring protection switching is now extended to Ethernet links. You can configure Ethernet ring protection for a series of two or more systems so that if one link fails, traffic is rerouted around the failure on the ring.

The basic idea of Ethernet ring protection is to use one specific link to protect the whole ring. This special link is the ring protection link (RPL). When all links are up and running, the RPL blocks traffic and remains idle. The RPL itself is controlled by the designated RPL owner node. There is only one RPL owner node on the ring and the RPL owner node is responsible for blocking the RPL interface under normal operating conditions. However, if a link failure occurs on the ring, the RPL owner node is responsible for unblocking the RPL interface and protection-switching the traffic on the alternate path around the ring. An Ethernet ring automatic protection switching (R-APS) messaging protocol coordinates the protection activities of all nodes on the ring. The APS blocks traffic over the failed link and unblocks traffic over the RPL.

When the failed link is repaired, the traffic reverts to its normal pattern. That is, the RPL owner blocks the RPL link and unblocks traffic over the cleared link.

Two or more nodes form a ring. Links between the nodes form a chain, with the last node also connecting the first. Every ring node therefore has two ports related to the ring, one in each direction. In this chapter, these directions are referred to as east and west.

Every node on the ring is one of two types:

- RPL owner node—This node owns the RPL and blocks or unblocks the RPL as conditions require. This node initiates the R-APS message.
- Normal node—All other nodes on the ring (that is, those that are not the RPL owner node) operate as normal nodes and have no special role on the ring.

In addition to roles, each node on the Ethernet ring can be in one of several states:

- **Init**—The node is not yet participating in the ring.
- **Idle**—The node is performing normally (there is no link failure on the ring). In this state, traffic is unblocked on both ring ports, except for the RPL owner node, which blocks the RPL port (the other RPL owner port is unblocked).
- **Protection**—When a failure occurs on the ring, a normal node will have traffic blocked on the ring port that connects to the failed link. The RPL owner, if it is not at one end of the failed link, will then unblock the RPL port so both ports are active.



NOTE: The R-APS protocol does not detect the number of RPL owner nodes configured on the ring. You must configure only one RPL and RPL owner per ring or protection switching will not work properly.

Ethernet ring protection only works when one link on the ring fails. Multiple link failures will break the ring and cause protection switching to fail.

Several restrictions apply to Ethernet ring protection:

- The Ethernet ring protection configured as a single instance only works at the physical level (adjacent nodes must be directly connected). The ring protection operates at the interface (port) level and not at the VLAN level.
- Manual (command-based) switching to protection mode is not supported.
- Nonrevertive switching is not supported. When the link failure is cleared, traffic always returns to normal operation.
- The interconnection of multiple rings for protection purposes is not supported.

You can configure Ethernet ring protection to optimize traffic load-balancing by using multiple ring instances. For more information about multiple ring instances, see [“Ethernet Ring Protection Using Ring Instances for Load Balancing” on page 15](#)

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Example: Configuring Ethernet Ring Protection for MX Series Routers on page 61](#)
- [Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 85](#)
- [Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 87](#)
- [Ethernet Ring Protection Using Ring Instances for Load Balancing on page 15](#)
- [Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers on page 68](#)

Ethernet Ring Protection Using Ring Instances for Load Balancing

Juniper Network MX Series 3D Universal Edge Routers support Ethernet ring protection (ERP) to help achieve high reliability and network stability. ERP is used in router or bridge networks to protect against link failure. A single-ring topology is configured that uses one specific link called a ring protection link (RPL) to protect the whole ring. When all links are up and running, the RPL blocks traffic and remains idle. However, if a link fails, the RPL routes traffic to bypass the failure on the ring.



NOTE: To learn how ERP works in a single-ring topology, see [“Ethernet Ring Protection” on page 13](#).

MX Series routers now support ERP ring instances. Whereas traffic in a single-ring topology follows the same path, traffic within ring instances allows some traffic to pass through one path while other traffic can follow a different path. Dividing traffic in this way supports traffic load balancing in the physical ring.

Ring instances are like traffic channels that contain different sets of virtual LANS (VLANs). A ring instance is responsible for the protection of a subset of VLANs that transport traffic over the physical ring. When ring instances are configured for the ring, each ring instance should have its own RPL owner, an east and a west interface, and a ring protection link end.

Each ring instance has a control channel and a specific data channel. A data channel is a group of bridge domain VLAN IDs. All VLAN IDs within the same ring interface share the same data-forwarding properties controlled by the ERP. If no data channel is defined in the ring configuration, ERP will only operate on the physical link instead of as a ring instance using logical links.

When operating ERP in a topology with other protocols, the following considerations should be observed:

- If a physical interface is part of an Ethernet ring, it cannot be configured for Spanning Tree Protocol (STP) or Multiple Spanning Tree Protocol (MSTP).
- ERP and Per-VLAN Spanning Tree (PVST) can be configured on the same topology as long as PVST doesn't share the same VLAN with any Ethernet ring instance configured on the physical port.
- If ERP is configured only as a physical ring instance (a ring without a data channel) in a topology also configured for PVST, ERP checks the PVST configuration on two ring interfaces and automatically creates a data channel excluding VLANs used by PVST.

Related Documentation

- [Ethernet Ring Protection on page 13](#)
- [Example: Configuring Ethernet Ring Protection for MX Series Routers on page 61](#)

PART 2

Configuration

- [ITU-T Y.1731 Ethernet Frame Delay Measurements Tasks on page 19](#)
- [IEEE 802.1ag OAM Connectivity-Fault Management Examples on page 25](#)
- [ITU-T Y.1731 Ethernet Frame Delay Measurements Examples on page 41](#)
- [IEEE 802.1ah OAM Link-Fault Management Examples on page 53](#)
- [Ethernet Ring Protection Examples on page 61](#)

CHAPTER 5

ITU-T Y.1731 Ethernet Frame Delay Measurements Tasks

- [Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements on page 19](#)
- [Triggering an Ethernet Frame Delay Measurements Session on page 20](#)
- [Viewing Ethernet Frame Delay Measurements Statistics on page 21](#)

Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging Service Level Agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

Before you can perform Ethernet frame delay measurements on MX Series routers, you must have done the following:

- Configured Ethernet OAM and CFM correctly
- Prepared the measurement between two compatibly configured MX Series routers
- Enabled the distributed periodic packet management daemon (ppmd)
- Avoided trying to perform Ethernet frame delay measurement on aggregated Ethernet or pseudowire interfaces, which are not supported
- Made sure the hardware-assisted timestamping is supported if that feature is configured

At the end of this configuration, you create two MX Series routers that can perform and display Ethernet frame delay measurements on Ethernet interfaces using optional hardware timestamping. By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

To configure hardware-assisted timestamping:

1. To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        performance-monitoring {
          hardware-assisted-timestamping; # Enable timestamping in hardware.
        }
      }
    }
  }
}
```

2. Ethernet frame delay measurement requires that distributed PPMD is enabled. Before you can gather statistics for Ethernet frame delay measurement, you must make sure that PPMD is configured properly. Without distributed PPMD, delay measurement results are not valid.

To perform Ethernet frame delay measurement, make sure that the following configuration statement is *NOT* present:

```
[edit routing-options]
ppm {
  no-delegate-processing; # This turns distributed PPMD OFF.
}
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Triggering an ETH-DM Session on page 20](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

Triggering an Ethernet Frame Delay Measurements Session

Before Ethernet frame delay measurement statistics can be displayed, they must be collected. To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement (one-way | two-way) (remote-mac-address | mep identifier) maintenance-domain name maintenance-association ma-id [count count] [wait time]** operational command.

The fields for this command are described in [Table 3 on page 21](#).

Table 3: Monitor Ethernet Delay Command Parameters

Parameter	Parameter Range	Description
one-way or two-way	NA	Perform a one-way or two-way (round-trip) delay measurement.
remote-mac-address	Unicast MAC address	Send delay measurement frames to the destination unicast MAC address (use the format xx:xx:xx:xx:xx:xx). Multicast MAC addresses are not supported.
mep identifier	1–8191	The MEP identifier to use for the measurement. The discovered MAC address for this MEP identifier is used.
maintenance-domain name	Existing MD name	Specifies an existing maintenance domain (MD) to use for the measurement.
maintenance-association ma-id	Existing MA identifier	Specifies an existing maintenance association (MA) identifier to use for the measurement.
count count	1–65535 (default: 10)	(Optional) Specifies the number of Ethernet frame delay frames to send. The default is 10.
wait time	1–255 seconds (default: 1)	(Optional) Specifies the number of seconds to wait between frames. The default is 1 second.

If you attempt to monitor delays to a nonexistent MAC address, you must exit the application manually using ^C:

```
user@host> monitor ethernet delay-measurement two-way 00:11:22:33:44:55
Two-way ETH-DM request to 00:11:22:33:44:55, Interface ge-5/2/9.0
^C
--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 0
Average delay: 0 usec, Average delay variation: 0 usec
Best case delay: 0 usec, Worst case delay: 0 usec
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

Viewing Ethernet Frame Delay Measurements Statistics

Once Ethernet frame delay measurement statistics have been collected, they can be displayed.

To retrieve the last 100 Ethernet frame delay measurement statistics per remote MEP or per CFM session, two types of **show** commands are provided:

- For all OAM frame counters and Ethernet frame delay measurement statistics
- For Ethernet frame delay measurement statistics only

To retrieve all Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.

To retrieve only Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.



NOTE: The only difference in the two commands is the use of the **mep-statistics** and **delay-statistics** keyword.

The fields for these commands are described in [Table 4 on page 22](#).

Table 4: Show Ethernet Delay Command Parameters

Parameter	Parameter Range	Description
maintenance-domain <i>name</i>	Existing MD name	Specifies an existing maintenance domain (MD) to use.
maintenance-association <i>ma-id</i>	Existing MA identifier	Specifies an existing maintenance association (MA) identifier to use.
local-mep <i>identifier</i>	1–8191	When a MEP has been specified, display statistics only for the local MEP.
remote-mep <i>identifier</i>	1–8191	When a MEP has been specified, display statistics only for the discovered MEP.
count <i>count</i>	1–100 (default:100)	The number of entries to display in the results table. By default, all 100 entries are displayed if they exist.



NOTE: For each MEP, you will see frame counters for sent and received Ethernet frame delay measurement frames whenever MEP statistics are displayed.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)

- [Triggering an ETH-DM Session on page 20](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

CHAPTER 6

IEEE 802.1ag OAM Connectivity-Fault Management Examples

- [Ethernet Operations, Administration, and Maintenance on page 25](#)
- [Ethernet OAM Connectivity Fault Management on page 26](#)
- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Ethernet Operations, Administration, and Maintenance

This topic provides an overview to help you effectively configure Ethernet Operations, Administration, and Maintenance (OAM) on a network of Juniper Networks MX Series 3D Universal Edge Routers. For more information about configuring OAM parameters on Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual local area network (VLAN) identifier for troubleshooting
- Work independently of the actual Ethernet transport and function over physical Ethernet ports, or a virtual service such as pseudowire, and so on.
- Isolate faults over a flat (or single operator) network architecture or a nested or hierarchical (or multi-provider) networks.

OAM can provide simple link-level information, provide performance statistics, or track end-to-end connectivity across the network. Simple link fault management (LFM) for Ethernet links is defined in IEEE 802.3ah.

Ethernet OAM functions are implemented as:

- Fault detection and notification (provided by continuity check messages)
- Path discovery (provided by the linktrace protocol)

- Fault isolation, verification, and recovery (isolation and verification are provided by a combination of protocols, while recovery is the function of protocols such as spanning tree)

The loopback protocol used in Ethernet OAM is modeled on the standard IP ping. After a fault is detected, the loopback protocol performs fault verification and isolation under the direction of a network operator. The loopback is performed using request and response message pairs. A unicast loopback message is generated by a MEP and a loopback reply is generated by the destination MIP or MEP. The target MAC address is learned by the continuity check protocol or linktrace protocol. The loopback message's packet is always forwarded to a unique port by the originating MEP, as determined by a MAC table lookup or the MEP interface MAC address. The target MIP or MEP generates a unicast loopback reply in response to the received loopback message. The loopback message follows the same path as a data packet, and intermediate bridges simply forward the packet to the destination MIP or MEP.

**Related
Documentation**

- *Ethernet OAM Feature Guide for MX Series Routers*
- [Ethernet OAM Connectivity Fault Management on page 4](#)
- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Ethernet OAM Connectivity Fault Management

The most complete connectivity fault management (CFM) is defined in IEEE 802.1ag. This topic emphasizes the use of CFM in a Metro Ethernet environment.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol which discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.
- Fault isolation using the loopback protocol. Similar to IP ping, this protocol works with the continuity check protocol during troubleshooting.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible. Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outermost domains are assigned a higher level than the innermost domains. Customer end points have to highest maintenance domain level. In a CFM

maintenance domain, each service instance is called a maintenance association. A maintenance association can be thought as a full mesh of maintenance endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

MEPs can be *up MEPs* or *down MEPs*. A link can connect a MEP at level 5 to a MEP at level 7. The interface at level 5 is an up MEP (because the other end of the link is at MEP level 7) and the interface at level 7 is a down MEP (because the other end of the link is at MEP level 5).

In a Metro Ethernet network, CFM is commonly used at two levels:

- By the service provider to check the connectivity among its provider edge (PE) routers
- By the customer to check the connectivity among its customer edge (CE) routers



NOTE: The configured customer CFM level must be greater than service provider CFM level.

In many Metro Ethernet networks, CFM is used to monitor connectivity over a VPLS and bridge network.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Operations, Administration, and Maintenance on page 3](#)
- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Example: Configuring Ethernet CFM over VPLS

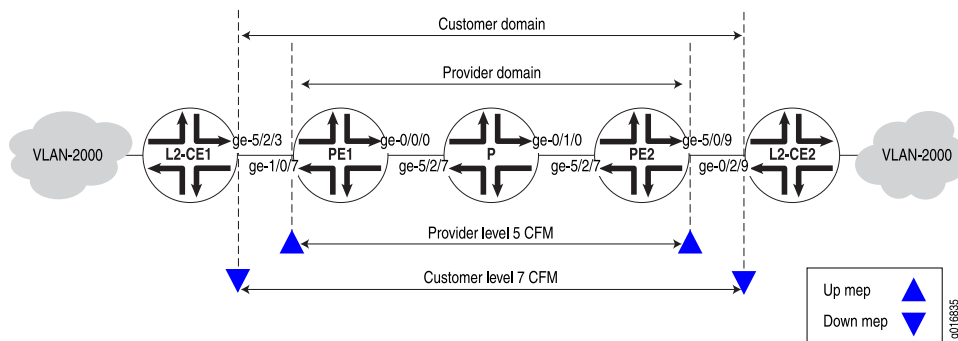
In this example, both the customer and service provider are running Ethernet CFM over a VPLS and a multiprotocol label switching (MPLS) network. The network is shown in [Figure 2 on page 28](#). The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1, P, and PE2.



NOTE: The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

The service provider is using CFM level 5 and the customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 2: Ethernet OAM with VPLS



The following are the configurations of the VPLS and CFM on the service provider routers.

Configuration of PE1

```
[edit chassis]
fpc 5 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}

[edit interfaces]
ge-1/0/7 {
  encapsulation flexible-ethernet-services;
  vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.200.1.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.231/32 {
        primary;
      }
      address 127.0.0.1/32;
    }
  }
}

[edit routing-instances]
vpls-vlan2000 {
```

```

instance-type vpls;
vlan-id 2000;
interface ge-1/0/7.1;
route-distinguisher 10.255.168.231:2000;
vrf-target target:1000:1;
protocols {
  vpls {
    site-range 10;
    site vlan2000-PE1 {
      site-identifier 2;
    }
  }
}

[edit protocols]
rsvp {
  interface ge-0/0/0.0;
}
mpls {
  label-switched-path PE1-to-PE2 {
    to 10.100.1.1;
  }
  interface ge-0/0/0.0;
}
bgp {
  group PE1-to-PE2 {
    type internal;
    local-address 10.200.1.1;
    family l2vpn {
      signaling;
    }
    local-as 65000;
    neighbor 10.100.1.1;
  }
}
ospf {
  traffic-engineering;
  reference-bandwidth 4g;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface ge-0/0/0.0;
  }
}
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain customer-site1 {
        level 5;
        maintenance-association customer-site1 {
          continuity-check {
            interval 1s;
          }
        }
      }
    }
  }
}

```

```
    }
    mep 100 {
      interface ge-1/0/7.1;
      direction up;
      auto-discovery;
    }
  }
}
}
```

Configuration of PE2

```
[edit chassis]
fpc 5 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}

[edit interfaces]
ge-5/0/9 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-5/2/7 {
  unit 0 {
    family inet {
      address 10.100.1.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.230/32 {
        primary;
      }
    }
    address 127.0.0.1/32;
  }
}

[edit routing-instances]
vpls-vlan2000 {
  instance-type vpls;
  vlan-id 2000;
  interface ge-5/0/9.1;
  route-distinguisher 10.255.168.230:2000;
```

```

vrf-target target:1000:1;
protocols {
  vpls {
    site-range 10;
    site vlan2000-PE2 {
      site-identifier 1;
    }
  }
}

[edit protocols]
rsvp {
  interface ge-5/2/7.0;
}
mpls {
  label-switched-path PE2-to-PE1 {
    to 10.200.1.1;
  }
  interface ge-5/2/7.0;
}
bgp {
  group PE2-to-PE1 {
    type internal;
    local-address 10.100.1.1;
    family l2vpn {
      signaling;
    }
    local-as 65000;
    neighbor 10.200.1.1;
  }
}
ospf {
  traffic-engineering;
  reference-bandwidth 4g;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface ge-5/2/7.0;
  }
}
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain customer-site1 {
        level 5;
        maintenance-association customer-site1 {
          continuity-check {
            interval 1s;
          }
        }
        mep 200 {
          interface ge-5/0/9.1;
          direction up;
        }
      }
    }
  }
}

```

```

        auto-discovery;
    }
}
}
}
}
}

```

Configuration of P router

MPLS only, no CFM needed:

```

[edit]
interfaces {
  ge-5/2/7 {
    # Connected to PE1
    unit 0 {
      family inet {
        address 10.200.1.10/24;
      }
      family mpls;
    }
  }
  ge-0/1/0 {
    # Connected to PE2
    unit 0 {
      family inet {
        address 10.100.1.10/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.168.240/32;
      }
    }
  }
}

```

```

[edit]
protocols {
  rsvp {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  mpls {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {

```

```

        disable;
    }
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
}
}
}

```

CFM on L2-CE1 Here is the configuration of CFM on L2-E1:

```

[edit interfaces]
ge-5/2/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 800 {
          interface ge-5/2/3.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
}
}

```

CFM on L2-CE2 Here is the configuration of CFM L2-CE2:

```

[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
      }
    }
  }
}

```

```

    }
    mep 700 {
        interface ge-0/2/9.0;
        direction down;
        auto-discovery;
    }
}
}
}
}
}

```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Operations, Administration, and Maintenance on page 3](#)
- [Ethernet OAM Connectivity Fault Management on page 4](#)
- [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Example: Configuring Ethernet CFM on Bridge Connections

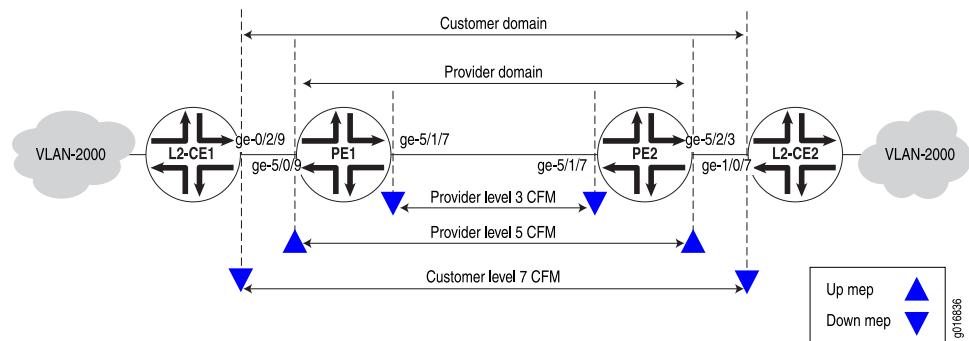
In this example, both the customer and service provider are running Ethernet CFM over a simple bridge network. The network is shown in [Figure 3 on page 34](#). The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1 and PE2.



NOTE: The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

The service provider is using CFM level 3 for the link between PE1 and PE2 and level 5 from one CE facing port to the other. The customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 3: Ethernet CFM over a Bridge Network



Here are the configurations of CFM on the customer routers.

CFM on L2-CE1

```
[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam ethernet]
connectivity-fault-management {
  maintenance-domain customer {
    level 7;
    maintenance-association customer-site1 {
      continuity-check {
        interval 1s;
      }
      mep 700 {
        interface ge-0/2/9.0;
        direction down;
        auto-discovery;
      }
    }
  }
}
```

CFM on L2-CE2

```
[edit interfaces]
ge-1/0/7 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam ethernet]
connectivity-fault-management {
  maintenance-domain customer {
    level 7;
    maintenance-association customer-site2 {
      continuity-check {
        interval 1s;
      }
      mep 800 {
        interface ge-1/0/7.0;
        direction down;
        auto-discovery;
      }
    }
  }
}
```

Here are the configurations of CFM on the provider routers.

CFM on PE1

```
[edit interfaces]
ge-5/0/9 {
  vlan-tagging;
```

```

encapsulation flexible-ethernet-services;
unit 0 {
    encapsulation vlan-bridge;
    vlan-id 2000;
}
}
ge-5/1/7 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}

[edit bridge-domains]
bridge-vlan2000 {
    domain-type bridge;
    vlan-id 2000;
    interface ge-5/0/9.0;
    interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
    level 5;
    maintenance-association provider-outer-site1 {
        continuity-check {
            interval 1s;
        }
        mep 200 {
            interface ge-5/0/9.0;
            direction up;
            auto-discovery;
        }
    }
}
maintenance-domain provider-inner {
    level 3;
    maintenance-association provider-inner-site1 {
        continuity-check {
            interval 1s;
        }
        mep 200 {
            interface ge-5/1/7.0;
            direction down;
            auto-discovery;
        }
    }
}

```

CFM on PE2

```

[edit interfaces]
ge-5/1/7 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
}

```

```

    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}
ge-5/2/3 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}

[edit bridge-domains]
bridge-vlan2000 {
    domain-type bridge;
    interface ge-5/2/3.0;
    interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
    level 5;
    maintenance-association provider-outer-site1 {
        continuity-check {
            interval 1s;
        }
        mep 100 {
            interface ge-5/2/3.0;
            direction up;
            auto-discovery;
        }
    }
}
maintenance-domain provider-inner {
    level 3;
    maintenance-association provider-inner-site1 {
        continuity-check {
            interval 1s;
        }
        mep 100 {
            interface ge-5/1/7.0;
            direction down;
            auto-discovery;
        }
    }
}

```

**Related
Documentation**

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Operations, Administration, and Maintenance on page 3](#)
- [Ethernet OAM Connectivity Fault Management on page 4](#)

- [Example: Configuring Ethernet CFM over VPLS on page 27](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces on page 38](#)

Example: Configuring Ethernet CFM on Physical Interfaces

CFM can be used to monitor the physical link between two routers. This functionality is similar to that supported by the IEEE 802.3ah LFM protocol.

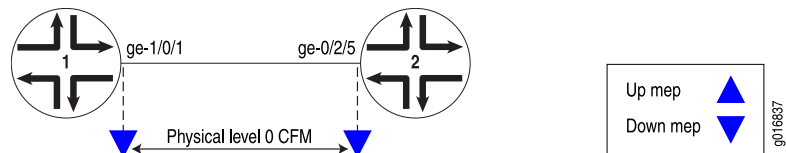
In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.



NOTE: The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

In the following example, two routers (Router 1 and Router 2) are connected by a point-to-point Gigabit Ethernet link. The link between these two routers is monitored using CFM. This is shown in [Figure 4 on page 38](#). The single boundary is a “down mep” in CFM terminology.

Figure 4: Ethernet CFM on Physical Interfaces



Router 1 Configure the interface and CFM:

```
[edit]
interfaces ge-1/0/1 {
  unit 0 {
    family inet;
  }
}

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
          }
        }
      }
    }
  }
}
```

```

        mep 100 {
            interface ge-1/0/1;
            direction down;
            auto-discovery;
        }
    }
}

```

The configuration on Router 2 mirrors that on Router 1, with the exception of the *mep-id*.

Router 2 Configure the interface and CFM:

```

[edit]
interfaces ge-0/2/5 {
    unit 0 {
        family inet;
    }
}

protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain private {
                    level 0;
                    maintenance-association private-ma {
                        continuity-check {
                            interval 1s;
                        }
                        mep 200 {
                            interface ge-0/2/5;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
    }
}

```

- Related Documentation**
- [Ethernet OAM Feature Guide for MX Series Routers](#)
 - [Ethernet Operations, Administration, and Maintenance on page 3](#)
 - [Ethernet OAM Connectivity Fault Management on page 4](#)
 - [Example: Configuring Ethernet CFM over VPLS on page 27](#)
 - [Example: Configuring Ethernet CFM on Bridge Connections on page 34](#)

CHAPTER 7

ITU-T Y.1731 Ethernet Frame Delay Measurements Examples

- Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces on page 41
- Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces on page 46
- Example: Configuring Ethernet Frame Delay Measurements with Untagged Interfaces on page 50

Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).



NOTE: These are not complete router configurations.

Configuration on Router **MX-1**:

```
[edit]
interfaces {
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
      }
    }
  }
}
```

```
linktrace {  
  path-database-size 255;  
  age 10s;  
}  
maintenance-domain md6 {  
  level 6;  
  maintenance-association ma6 {  
    continuity-check {  
      interval 100ms;  
      hold-interval 1;  
    }  
    mep 201 {  
      interface ge-5/2/9.0;  
      direction down;  
      auto-discovery;  
    }  
  }  
}  
}  
}  
}
```

Configuration on Router **MX-2**:

```
[edit]  
interfaces {  
  ge-0/2/5 {  
    vlan-tagging;  
    unit 0 {  
      vlan-id 512;  
    }  
  }  
}  
protocols {  
  oam {  
    ethernet {  
      connectivity-fault-management {  
        traceoptions {  
          file eoam_cfm.log size 1g files 2 world-readable;  
          flag all;  
        }  
        linktrace {  
          path-database-size 255;  
          age 10s;  
        }  
        maintenance-domain md6 {  
          level 6;  
          maintenance-association ma6 {  
            continuity-check {  
              interval 100ms;  
              hold-interval 1;  
            }  
            mep 101 {  
              interface ge-0/2/5.0;  
              direction down;  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

```

        auto-discovery;
    }
}
}
}
}
}
}
}

```

From Router **MX-2**, start a one-way delay measurement to Router **MX-1**.

```

user@MX-2> monitor ethernet delay-measurement one-way mep 201 maintenance-domain md6
maintenance-association ma6 count 10
One-way ETH-DM request to 00:90:69:0a:43:94, Interface ge-0/2/5.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA

```

The counters are displayed as part of the local MEP database on Router **MX-2**.

```

user@MX-2> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-association ma6
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMS sent                                  : 1590
  CCMS received out of sequence              : 0
  LBMS sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 10
  Valid 1DMs received                       : 0
  Invalid 1DMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                      : 0
  Invalid DMRs received                    : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    201      00:90:69:0a:43:94  ok    ge-0/2/5.0

```

The remote MEP database statistics are available on Router **MX-1**.

```

user@MX-1> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
Auto-discovery: enabled, Priority: 0
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1572
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received          : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 0
  Valid 1DMs received                       : 10
  Invalid 1DMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                      : 0
  Invalid DMRs received                    : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    101      00:90:69:0a:48:57  ok    ge-5/2/9.0

```

The remote Router **MX-1** should also collect the delay statistics (up to 100 per session) for display with **mep-statistics** or **delay-statistics**.

```

user@MX-1> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md6
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
  CCMs sent                                  : 3240
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received          : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 0
  Valid 1DMs received                       : 10
  Invalid 1DMs received                     : 0
  DMMs sent                                  : 0

```

```

DMRs sent : 0
Valid DMRs received : 0
Invalid DMRs received : 0

```

```

Remote MEP identifier: 101
Remote MAC address: 00:90:69:0a:48:57

```

```

Delay measurement statistics:
Index One-way delay Two-way delay
      (usec)      (usec)

```

```

1      370
2      357
3      344
4      332
5      319
6      306
7      294
8      281
9      269
10     255

```

```

Average one-way delay : 312 usec
Average one-way delay variation: 11 usec
Best case one-way delay : 255 usec
Worst case one-way delay : 370 usec

```

```

user@MX-1> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6

```

```

MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1

```

```

Remote MAC address: 00:90:69:0a:48:57

```

```

Delay measurement statistics:
Index One-way delay Two-way delay
      (usec)      (usec)

```

```

1      370
2      357
3      344
4      332
5      319
6      306
7      294
8      281
9      269
10     255

```

```

Average one-way delay : 312 usec
Average one-way delay variation: 11 usec
Best case one-way delay : 255 usec

```



NOTE: When two systems are close to each other, their one-way delay values are very high compared to their two-way delay values. This is because one-way delay measurement requires the timing for the two systems to be synchronized at a very granular level and MX Series routers do not support this granular synchronization. However, two-way delay measurement does not require synchronized timing, making two-way delay measurements more accurate.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
- [Triggering an ETH-DM Session on page 20](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).



NOTE: These are not complete router configurations.

Configuration on Router **MX-1**:

```
[edit]
interfaces {
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
        maintenance-domain md6 {
          level 6;
          maintenance-association ma6 {
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
          }
        }
      }
    }
  }
}
```

```

        mep 201 {
            interface ge-5/2/9.0;
            direction down;
            auto-discovery;
        }
    }
}

```

Configuration on Router **MX-2**:

```

[edit]
interfaces {
    ge-0/2/5 {
        vlan-tagging;
        unit 0 {
            vlan-id 512;
        }
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
                linktrace {
                    path-database-size 255;
                    age 10s;
                }
                maintenance-domain md6 {
                    level 6;
                    maintenance-association ma6 {
                        continuity-check {
                            interval 100ms;
                            hold-interval 1;
                        }
                        mep 101 {
                            interface ge-0/2/5.0;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
    }
}

```

From Router **MX-1**, start a two-way delay measurement to Router **MX-2**.

```
user@MX-1> monitor ethernet delay-measurement two-way mep 101 maintenance-domain md6
maintenance-association ma6 count 10
```

```
Two-way ETH-DM request to 00:90:69:0a:48:57, Interface ge-5/2/9.0
DMR received from 00:90:69:0a:48:57 Delay: 100 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 8 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 111 usec Delay variation: 19 usec
DMR received from 00:90:69:0a:48:57 Delay: 110 usec Delay variation: 1 usec
DMR received from 00:90:69:0a:48:57 Delay: 119 usec Delay variation: 9 usec
DMR received from 00:90:69:0a:48:57 Delay: 122 usec Delay variation: 3 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 30 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 108 usec Delay variation: 16 usec
```

```
--- Delay measurement statistics ---
```

```
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

The counters are displayed as part of the MEP database on Router **MX-1** maintenance domain **MD6**.

```
user@MX-1> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6
```

```
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
Auto-discovery: enabled, Priority: 0
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no
  Cross-connect CCM received                   : no
  RDI sent by some MEP                         : no
Statistics:
  CCMS sent                                   : 894
  CCMS received out of sequence                : 0
  LBMS sent                                   : 0
  Valid in-order LBRs received                 : 0
  Valid out-of-order LBRs received             : 0
  LBRs received with corrupted data            : 0
  LBRs sent                                   : 0
  LTMs sent                                   : 0
  LTMs received                               : 0
  LTRs sent                                   : 0
  LTRs received                               : 0
  Sequence number of next LTM request          : 0
  1DMs sent                                   : 0
  Valid 1DMs received                         : 0
  Invalid 1DMs received                       : 0
  DMMs sent                                   : 10
  DMRs sent                                   : 0
  Valid DMRs received                         : 10
  Invalid DMRs received                       : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    101      00:90:69:0a:48:57  ok    ge-5/2/9.0
```

The collected MEP statistics are saved (up to 100 per remote MEP or per CFM session) and displayed as part of the MEP statistics on Router **MX-1**.

```
user@MX-1> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md6
```

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
```

```
Remote MEP count: 1
```

```
CCMs sent : 3154
CCMs received out of sequence : 0
LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
1DMs sent : 0
Valid 1DMs received : 0
Invalid 1DMs received : 0
DMMs sent : 10
DMRs sent : 0
Valid DMRs received : 10
Invalid DMRs received : 0
```

```
Remote MEP identifier: 101
```

```
Remote MAC address: 00:90:69:0a:48:57
```

```
Delay measurement statistics:
```

Index	One-way delay (usec)	Two-way delay (usec)
1		100
2		92
3		92
4		111
5		110
6		119
7		122
8		92
9		92
10		108

```
Average two-way delay : 103 usec
```

```
Average two-way delay variation: 8 usec
```

```
Best case two-way delay : 92 usec
```

```
Worst case two-way delay : 122 usec
```

The collected delay statistics are also saved (up to 100 per session) and displayed as part of the MEP delay statistics on Router **MX-1**.

```
user@MX-1> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6
```

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
```

```
Remote MEP count: 1
```

```
Remote MAC address: 00:90:69:0a:48:57
```

```
Delay measurement statistics:
```

Index	One-way delay (usec)	Two-way delay (usec)
-------	-------------------------	-------------------------

1	100
2	92
3	92
4	111
5	110
6	119
7	122
8	92
9	92
10	108
Average two-way delay : 103 usec	
Average two-way delay variation: 8 usec	
Best case two-way delay : 92 usec	
Worst case two-way delay : 122 usec	

- Related Documentation**
- [Ethernet OAM Feature Guide for MX Series Routers](#)
 - [Ethernet Frame Delay Measurements on page 7](#)
 - [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
 - [Triggering an ETH-DM Session on page 20](#)
 - [Viewing ETH-DM Statistics on page 21](#)
 - [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
 - [Configuring ETH-DM with Untagged Interfaces on page 50](#)

Example: Configuring Ethernet Frame Delay Measurements with Untagged Interfaces

Ethernet frame delay measurements are supported on untagged interfaces. All commands are the same as for tagged interfaces. Only the configurations are different. This section shows the untagged interface configurations for Routers **MX-1** and **MX-2**.



NOTE: These are not complete router configurations.

Untagged interface configuration for Router **MX-1**.

```
[edit]
interfaces {
  ge-5/0/0 {
    unit 0;
  }
  ge-5/2/9 {
    unit 0;
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
      }
    }
  }
}
```

```

    }
    linktrace {
        path-database-size 255;
        age 10s;
    }
    maintenance-domain md6 {
        level 6;
        maintenance-association ma6 {
            continuity-check {
                interval 100ms;
                hold-interval 1;
            }
            mep 201 {
                interface ge-5/0/0;
                direction down;
                auto-discovery;
            }
        }
    }
}
}
}
}
}
}
}

```

Untagged interface configuration for Router **MX-2**.

```

[edit]
interfaces {
    ge-0/2/2 {
        unit 0;
    }
    ge-0/2/5 {
        unit 0;
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
            }
            linktrace {
                path-database-size 255;
                age 10s;
            }
            maintenance-domain md6 {
                level 6;
                maintenance-association ma6 {
                    continuity-check {
                        interval 100ms;
                        hold-interval 1;
                    }
                    mep 101 {
                        interface ge-0/2/2;

```

```
        direction down;  
        auto-discovery;  
    }  
    }  
    }  
    }  
    }  
}
```

**Related
Documentation**

- *Ethernet OAM Feature Guide for MX Series Routers*
- [Ethernet Frame Delay Measurements on page 7](#)
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
- [Triggering an ETH-DM Session on page 20](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)

IEEE 802.1ah OAM Link-Fault Management Examples

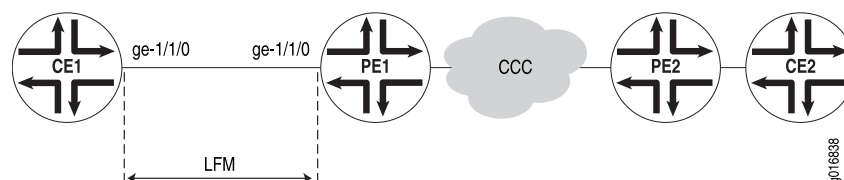
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 53](#)
- [Example: Configuring Ethernet LFM for CCC on page 54](#)
- [Example: Configuring Ethernet LFM for Aggregated Ethernet on page 56](#)
- [Example: Configuring Ethernet LFM with Loopback Support on page 58](#)

Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge

In this example, LFM is enabled on an IP link between the provider edge (PE) and customer edge (CE) interfaces. If the link goes down, the fault will be detected by LFM and the interfaces on both sides will be marked **Link-Layer-Down**. This results in notifications to various subsystems (for example, routing) which will take appropriate action.

The link running LFM is shown in [Figure 5 on page 53](#).

Figure 5: Ethernet LFM Between Provider Edge and Customer Edge



To configure Ethernet LFM on an IP link between PE and CE interfaces:

1. Configure LFM on the PE router:

```
[edit]
interfaces ge-1/1/0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {
```

```
    ethernet {  
        link-fault-management {  
            interface ge-1/1/0 {  
                pdu-interval 1000;  
                pdu-threshold 5;  
            }  
        }  
    }  
}
```

2. Configure LFM on the CE router:

```
[edit]  
interfaces ge-1/1/0 {  
    unit 0 {  
        family inet {  
            address 11.11.11.2/24;  
        }  
    }  
}  
protocols {  
    oam {  
        ethernet {  
            link-fault-management {  
                interface ge-1/1/0 {  
                    pdu-interval 1000;  
                    pdu-threshold 5;  
                }  
            }  
        }  
    }  
}
```

**Related
Documentation**

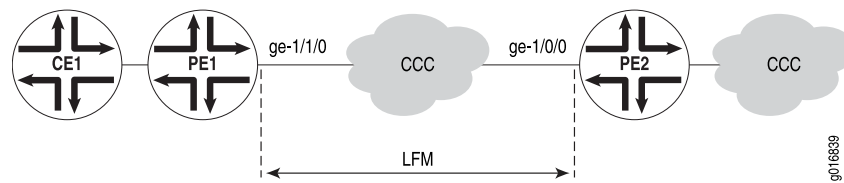
- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet OAM Link Fault Management on page 11](#)
- [Example: Configuring Ethernet LFM for CCC on page 54](#)
- [Example: Configuring Ethernet LFM for Aggregated Ethernet on page 56](#)
- [Example: Configuring Ethernet LFM with Loopback Support on page 58](#)

Example: Configuring Ethernet LFM for CCC

In this example, LFM is configured between two PEs (PE1 and PE2) connected using CCC. With LFM in place, a link fault will be detected immediately, instead of depending on routing protocols to find the fault on end-to-end CCC connection. This also helps in detecting the exact failed link instead of only finding that the end-to-end CCC connectivity has failed. Also, because LFM runs at the link-layer level, it does not need a IP address to operate and so can be used where bidirectional fault detection (BFD) cannot.

The links running LFM are shown in [Figure 6 on page 55](#)

Figure 6: Ethernet LFM for CCC



To configure Ethernet LFM between two PEs connected using CCC:

1. Configure LFM on the PE1 router with CCC:

```
[edit]
interfaces ge-1/1/0 {
  encapsulation ethernet-ccc;
  unit 0;
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/1/0 {
          pdu-interval 1000;
          pdu-threshold 5;
        }
      }
    }
  }
}
```

2. Configure LFM on the PE2 router with CCC:

```
[edit]
interfaces ge-1/0/0 {
  encapsulation ethernet-ccc;
  unit 0;
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/0/0 {
          pdu-interval 1000;
          pdu-threshold 5;
        }
      }
    }
  }
}
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet OAM Link Fault Management on page 11](#)
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 53](#)

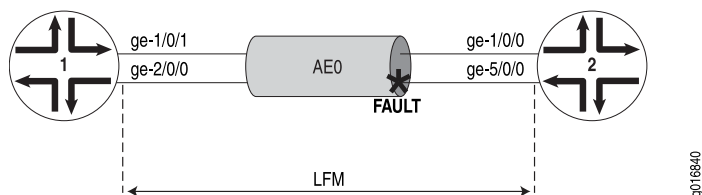
- [Example: Configuring Ethernet LFM for Aggregated Ethernet on page 56](#)
- [Example: Configuring Ethernet LFM with Loopback Support on page 58](#)

Example: Configuring Ethernet LFM for Aggregated Ethernet

In this example, LFM is configured on an aggregated Ethernet interface (AE0) between Router 1 and Router 2. When configured on aggregated Ethernet, LFM runs on all the individual member links. LFM is enabled or disabled on the member links as they are added or deleted from the aggregation group. The status of individual links is used to determine the status of the aggregated interface.

The use of LFM with aggregated Ethernet is shown in [Figure 7 on page 56](#).

Figure 7: Ethernet LFM for Aggregated Ethernet



To configure LFM on an aggregated Ethernet interface between two routers:

1. Configure LFM on Router 1 for AE0:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
interfaces ge-1/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ge-2/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ae0 {
  unit 0 {
    family inet {
      address 11.11.11.2/24;
    }
  }
}
protocols {
  oam {
```

```

    ethernet {
      link-fault-management {
        interface ae0;
      }
    }
  }
}

```

2. Configure LFM on Router 2 for AE0:

```

[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
interfaces ge-1/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ge-5/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ae0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ae0;
      }
    }
  }
}
}

```

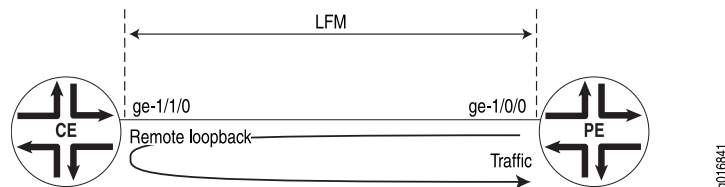
**Related
Documentation**

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet OAM Link Fault Management on page 11](#)
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 53](#)
- [Example: Configuring Ethernet LFM for CCC on page 54](#)
- [Example: Configuring Ethernet LFM with Loopback Support on page 58](#)

Example: Configuring Ethernet LFM with Loopback Support

In this example, LFM is configured between provider edge (PE) router and the customer edge (CE) router. The PE router can put the CE router in remote loopback mode. This allows the PE to have all the traffic sent to the CE router looped back for diagnostics purposes, as shown in [Figure 8 on page 58](#).

Figure 8: Ethernet LFM with Loopback Support



To configure LFM between a PE router and a CE router:

1. Configure LFM loopback on the PE router:

```
[edit]
interfaces ge-1/0/0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/0/0 {
          pdu-interval 1000;
          pdu-threshold 5;
          remote-loopback;
        }
      }
    }
  }
}
```

2. Configure LFM loopback on the CE router:

```
[edit]
interfaces ge-1/1/0 {
  unit 0 {
    family inet {
      address 11.11.11.2/24;
    }
  }
}
protocols {
  oam {
    ethernet {
```

```

link-fault-management {
  interface ge-1/1/0 {
    pdu-interval 1000;
    pdu-threshold 5;
    negotiation-options {
      allow-remote-loopback;
    }
  }
}

```



NOTE: If the negotiation options `allow-remote-loopback` statement on the CE router is deleted before removing the CE router from remote loopback mode, traffic flow between the PE router and CE router is affected. Hence, delete the `remote-loopback` statement on the PE router before deleting the `negotiation-options allow-remote-loopback` statement on the CE router.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet OAM Link Fault Management on page 11](#)
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 53](#)
- [Example: Configuring Ethernet LFM for CCC on page 54](#)
- [Example: Configuring Ethernet LFM for Aggregated Ethernet on page 56](#)

CHAPTER 9

Ethernet Ring Protection Examples

- [Example: Configuring Ethernet Ring Protection for MX Series Routers on page 61](#)
- [Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers on page 68](#)
- [Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 85](#)
- [Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 87](#)

Example: Configuring Ethernet Ring Protection for MX Series Routers

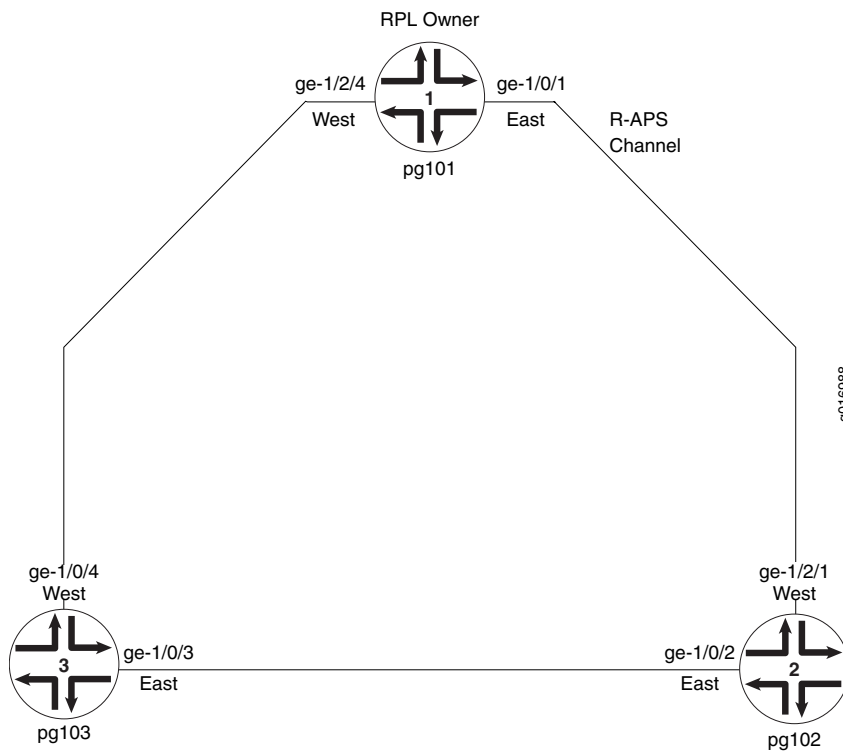
This example configures Ethernet ring protection for three MX Series router nodes:

- [Example Topology on page 61](#)
- [Router 1 \(RPL Owner\) Configuration on page 63](#)
- [Router 2 Configuration on page 64](#)
- [Router 3 Configuration on page 66](#)

Example Topology

The links connecting the three MX Series routers are shown in [Figure 9 on page 62](#).

Figure 9: Ethernet Ring Protection Example Nodes



This example uses the following topology details for Ethernet ring protection:

- Router 1 is the RPL owner. The node identification for Router 1 is MAC address **00:01:01:00:00:01**.
- The RPL link is **ge-1/0/1.1** (this is also the R-APS messaging control channel).
- Traffic flows among the nodes in the configured bridge domains. (That is, only the control channels are configured.)
- Router 1's east control channel interface is **ge-1/0/1.1** (the RPL) and the west control channel interface is **ge-1/2/4.1**. The protection group is **pg101**.
- Router 2's east control channel interface is **ge-1/0/2.1** (the RPL) and the west control channel interface is **ge-1/2/1.1**. The protection group is **pg102**.
- Router 3's east control channel interface is **ge-1/0/3.1** (the RPL) and the west control channel interface is **ge-1/0/4.1**. The protection group is **pg103**.



NOTE: Although not strictly required for physical ring protection, this example configures Ethernet OAM with MEPS.

Router 1 (RPL Owner) Configuration

To configure Router 1 (the RPL owner):

1. Configure the interfaces:

```
[edit]
interfaces {
  ge-1/0/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/2/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  irb {
    unit 0 {
      family inet {
        address address 192.1.1.1/24;
      }
    }
  }
}
```

2. Configure the bridge domain:

```
[edit]
bridge-domains {
  bd1 {
    domain-type bridge;
    vlan-id 100;
    interface ge-1/2/4.1;
    interface ge-1/0/1.1;
    routing-interface irb.0;
  }
}
```

3. Configure the Ethernet ring protection group:

```
[edit]
protocols {
  protection-group {
    ethernet-ring pg101 {
      node-id 00:01:01:00:00:01;
      ring-protection-link-owner;
      east-interface {
        control-channel ge-1/0/1.1;
      }
    }
  }
}
```

```
        ring-protection-link-end;
        west-interface {
            control-channel ge-1/2/4.1;
        }
    }
}
```

4. Configure Ethernet OAM:

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
      }
      maintenance-domain d1 {
        level 0;
        maintenance-association 100 {
          mep 1 {
            interface ge-1/0/1;
            remote-mep 2 {
              action-profile rmep-defaults;
            }
          }
        }
      }
      maintenance-domain d2 {
        level 0;
        maintenance-association 100 {
          mep 1 {
            interface ge-1/2/4;
            remote-mep 2 {
              action-profile rmep-defaults;
            }
          }
        }
      }
    }
  }
}
```

Router 2 Configuration

To configure Router 2:

1. Configure the interfaces:

```
[edit]
interfaces {
  ge-1/0/2 {
```

```

    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
        encapsulation vlan-bridge;
        vlan-id 100;
    }
}
ge-1/2/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
        encapsulation vlan-bridge;
        vlan-id 100;
    }
}
irb {
    unit 0 {
        family inet {
            address address 192.1.1.22/24;
        }
    }
}
}

```

2. Configure the bridge domain:

```

[edit]
bridge-domains {
    bd1 {
        domain-type bridge;
        vlan-id 100;
        interface ge-1/2/1.1;
        interface ge-1/0/2.1;
        routing-interface irb.0;
    }
}

```

3. Configure the Ethernet protection group:

```

[edit]
protocols {
    protection-group {
        ethernet-ring pg102 {
            node-id 00:22:22:22:22:22;
            east-interface {
                control-channel ge-1/0/2.1;
            }
            west-interface {
                control-channel ge-1/2/1.1;
            }
        }
    }
}

```

4. Configure Ethernet OAM:

```

[edit]
protocols {

```

```

oam {
  ethernet {
    connectivity-fault-management {
      action-profile rmep-defaults {
        default-action {
          interface-down;
        }
      }
    }
    maintenance-domain d1 {
      level 0;
      maintenance-association 100 {
        mep 2 {
          interface ge-1/2/1;
          remote-mep 1 {
            action-profile rmep-defaults;
          }
        }
      }
    }
    maintenance-domain d3 {
      level 0;
      maintenance-association 100 {
        mep 1 {
          interface ge-1/0/2;
          remote-mep 2 {
            action-profile rmep-defaults;
          }
        }
      }
    }
  }
}

```

Router 3 Configuration

To configure Router 3:

1. Configure the interfaces:

```

[edit]
interfaces {
  ge-1/0/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/0/3 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
    }
  }
}

```

```

        vlan-id 100;
    }
}
irb {
    unit 0 {
        family inet {
            address 192.1.1.33/24;
        }
    }
}
}

```

2. Configure the bridge domain:

```

[edit]
bridge-domains {
    bd1 {
        domain-type bridge;
        vlan-id 100;
        interface ge-1/0/4.1;
        interface ge-1/0/3.1;
        routing-interface irb.0;
    }
}

```

3. Configure the Ethernet protection group:

```

[edit]
protocols {
    protection-group {
        ethernet-ring pg103 {
            node-id 00:33:33:33:33:33;
            east-interface {
                control-channel ge-1/0/3.1;
            }
            west-interface {
                control-channel ge-1/0/4.1;
            }
        }
    }
}

```

4. Configure Ethernet OAM:

```

[edit]
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
            }
            maintenance-domain d2 {
                level 0;
                maintenance-association 100 {
                    mep 2 {

```

```
        interface ge-1/0/4;
        remote-mep 1 {
            action-profile rmep-defaults;
        }
    }
}
maintenance-domain d3 {
    level 0;
    maintenance-association 100 {
        mep 2 {
            interface ge-1/0/3;
            remote-mep 1 {
                action-profile rmep-defaults;
            }
        }
    }
}
}
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Ring Protection on page 13](#)
- [Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 85](#)
- [Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 87](#)

Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers

MX Series routers support Ethernet ring protection (ERP) to help achieve high reliability and network stability. ERP is used in router or bridge networks to protect against link failure. A single-ring topology is configured that uses one specific link called a ring protection link (RPL) to protect the whole ring. When all links are up and running, the RPL blocks traffic and remains idle. However, if a link fails, the RPL routes traffic to bypass the failure on the ring.

MX Series routers now support ERP ring instances. Whereas traffic in a ring topology follows the same path, traffic within a ring instance uses data channels to allow some traffic to pass through one path while other traffic can follow a different one. Dividing traffic in this way supports traffic load-balancing in the ring.

This example describes how to use ERP with ring instances to load-balance traffic while still providing network protection from link failure:

- [Requirements on page 69](#)
- [Overview and Topology on page 69](#)

- [Configuration on page 72](#)
- [Verification on page 80](#)

Requirements

This example uses the following hardware and software components:

- Two MX Series routers acting as core switches
- One MX Series router acting as an aggregation switch
- Junos OS Release 10.2 or later for MX Series routers

Overview and Topology

[Figure 10 on page 70](#) displays the topology for this example. The topology contains three MX Series routers. CS1 and CS2 act as core routers in the topology, and AS1 acts as an aggregation switch. Each router has two ring instances, ring-1 and ring-2. All nodes on the ring coordinate protection activities by exchanging messages through the Ethernet ring automatic protection switching (R-APS) messaging protocol. Each ring instance has an RPL owner. The ring-1 RPL owner is CS1; the ring-2 RPL owner is CS2. The RPL owners block or unblock the RPL as conditions require and initiate R-APS messages.

Each ring instance has two interface ports (an east interface and a west interface) that participate in the instance. Interface **ge-2/0/8.0**, the west interface on CS2, is the ring protection link end where ring-2's RPL terminates. Interface **ge-3/2/4.0**, the east interface on CS1, is the ring protection link end where ring-1's RPL terminates.

Each ring instance has a data channel. A data channel is a group of bridge domain virtual LAN (VLAN) IDs. All VLAN IDs within the same ring interface share the same data-forwarding properties controlled by the ERP. The data channel on ring-1 is [200, 300]. The data channel on ring-2 is [500, 600].

Two customer site switches are connected to AS1. Customer site 1 uses VLANs 200 and 300. Customer site 2 uses VLANs 500 and 600.

Figure 10: ERP with Multiple Protection Instances Configured on Three MX Series Routers

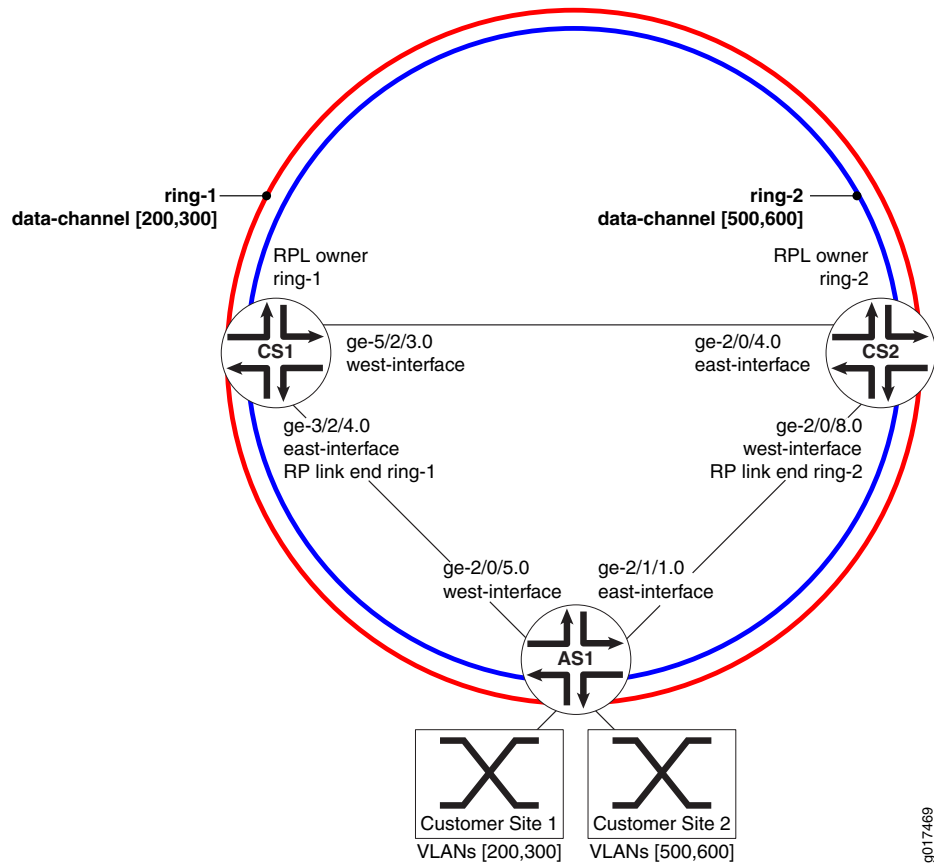


Table 5 on page 70 describes the components of the example topology.

Table 5: Components of the Network Topology

Property	Settings
Ring instances	<ul style="list-style-type: none"> ring-1—Data channel [200,300] ring-2—Data channel [500,600]
Customer sites	Two customer sites are connected to AS 1: <ul style="list-style-type: none"> Customer site 1, VLAN 200 and VLAN 300 Customer site 2, VLAN 500 and VLAN 600

Table 5: Components of the Network Topology (*continued*)

Property	Settings
CS1 router	<p>CS1 has the following protection group properties:</p> <ul style="list-style-type: none"> • RPL owner—ring-1. • East interface—ge-3/2/4.0. • West interface—ge-5/2/3.0. • Data channel for ring-1—VLAN 200, VLAN 300. • Data channel for ring-2—VLAN 500, VLAN 600. • Ring protection link end for ring-1—ge-3/2/4.0. <p>CS1 has the following routing and bridging properties:</p> <ul style="list-style-type: none"> • Routing instance—vs. • Bridge domains: <ul style="list-style-type: none"> • bd100 is associated with vlan-id 100. • bd101 is associated with vlan-id 101. • bd200 is associated with vlan-id 200. • bd300 is associated with vlan-id 300. • bd500 is associated with vlan-id 500. • bd600 is associated with vlan-id 600.
CS2 router	<p>CS2 has the following protection group properties:</p> <ul style="list-style-type: none"> • RPL owner—ring-2. • East interface—ge-2/0/4.0. • West interface—ge-2/0/8.0. • Ring protection link end for ring-2—ge-2/0/8.0. • Data channel for ring-1—VLAN 200, VLAN 300. • Data channel for ring-2—VLAN 500, VLAN 600. <p>CS2 has the following bridging properties:</p> <ul style="list-style-type: none"> • bd100 is associated with vlan-id 100. • bd101 is associated with vlan-id 101. • bd200 is associated with vlan-id 200. • bd300 is associated with vlan-id 300. • bd500 is associated with vlan-id 500. • bd600 is associated with vlan-id 600.

Table 5: Components of the Network Topology (*continued*)

Property	Settings
AS1 router	<p>AS1 has the following protection group properties:</p> <ul style="list-style-type: none"> • East interface—ge-2/0/5.0. • West interface—ge-2/1/1.0. • Data channel for ring-1—VLAN 200, VLAN 300. • Data channel for ring-2—VLAN 500, VLAN 600. <p>AS1 has the following bridging properties:</p> <ul style="list-style-type: none"> • bd100 is associated with vlan-id 100. • bd101 is associated with vlan-id 101. • bd200 is associated with vlan-id 200. • bd300 is associated with vlan-id 300. • bd500 is associated with vlan-id 500. • bd600 is associated with vlan-id 600.

Configuration

To enable ERP with ring instances on CS1, CS2, and AS1, perform these tasks:

- [Configuring ERP on CS1 on page 72](#)
- [Configuring ERP on CS2 on page 75](#)
- [Configuring ERP on AS1 on page 78](#)

Configuring ERP on CS1

CLI Quick Configuration

To quickly configure CS1 for ERP, copy the following commands and paste them into the switch terminal window of CS1:

```
[edit]
set interfaces ge-3/2/4 vlan-tagging
set interfaces ge-3/2/4 unit 0 family bridge interface-mode trunk
set interfaces ge-3/2/4 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-5/2/3 vlan-tagging
set interfaces ge-5/2/3 unit 0 family bridge interface-mode trunk
set interfaces ge-5/2/3 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 ring-protection-link-owner
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-3/2/4.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 east-interface ring-protection-link-end
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-5/2/3.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-3/2/4.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-5/2/3.0
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set routing-instances vs instance-type virtual-switch
set routing-instances vs interface ge-3/2/4.0
set routing-instances vs interface ge-5/2/3.0
set routing-instances vs bridge-domains bd101 vlan-id 101
```

```

set routing-instances vs bridge-domains bd200 vlan-id 200
set routing-instances vs bridge-domains bd300 vlan-id 300
set routing-instances vs bridge-domains bd500 vlan-id 500
set routing-instances vs bridge-domains bd600 vlan-id 600

```

Step-by-Step Procedure To configure ERP on CS1:

1. Configure the trunk interface **ge-3/2/4** to connect CS1 to CS2 and the trunk interface **ge-5/2/3** to connect CS1 to AS, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```

[edit interfaces]
user@cs1# set ge-3/2/4 vlan-tagging
user@cs1# set ge-3/2/4 unit 0 family bridge interface-mode trunk
user@cs1# set ge-3/2/4 unit 0 family bridge vlan-id-list 100-1000
user@cs1# set ge-5/2/3 vlan-tagging
user@cs1# set ge-5/2/3 unit 0 family bridge interface-mode trunk
user@cs1# set ge-5/2/3 unit 0 family bridge vlan-id-list 100-1000

```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**, and configure **ring-1** as the ring protection link owner:



NOTE: Always configure the east-interface statement first, before configuring the west-interface statement.

```

[edit protection-group]
user@cs1# set ethernet-ring ring-1 ring-protection-link-owner
user@cs1# set ethernet-ring ring-1 east-interface control-channel ge-3/2/4.0
user@cs1# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@cs1# set ethernet-ring ring-1 east-interface ring-protection-link-end
user@cs1# set ethernet-ring ring-1 west-interface control-channel ge-5/2/3.0
user@cs1# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@cs1# set ethernet-ring ring-1 data-channel vlan [200, 300]
user@cs1# set ethernet-ring ring-2 east-interface control-channel ge-3/2/4.0
user@cs1# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@cs1# set ethernet-ring ring-2 west-interface control-channel ge-5/2/3.0
user@cs1# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@cs1# set ethernet-ring ring-2 data-channel vlan [500, 600]

```

3. Configure the routing instance, the bridge domains, and the VLAN IDs associated with each bridge domain:

```

[edit routing-instances]
user@cs1# set vs instance-type virtual-switch
user@cs1# set vs interface ge-3/2/4.0
user@cs1# set vs interface ge-5/2/3.0
user@cs1# set vs bridge-domains bd100 vlan-id 100
user@cs1# set vs bridge-domains bd101 vlan-id 101
user@cs1# set vs bridge-domains bd200 vlan-id 200
user@cs1# set vs bridge-domains bd300 vlan-id 300
user@cs1# set vs bridge-domains bd500 vlan-id 500
user@cs1# set vs bridge-domains bd600 vlan-id 600

```

Results Check the results of the configuration:

```

user@cs1> show configuration
interfaces {
  ge-3/2/4 {

```

```
vlan-tagging;
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
    }
}
ge-5/2/3 {
    vlan-tagging;
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 100-1000;
        }
    }
}
protocols {
    protection-group {
        ethernet-ring ring-1 {
            east-interface {
                control-channel {
                    ge-3/2/4.0;
                    vlan 100;
                }
            }
            ring-protection-link-end;
        }
        west-interface {
            control-channel {
                ge-5/2/3.0;
                vlan 100;
            }
        }
        data-channel {
            vlan [ 200 300 ];
        }
    }
}
    protection-group {
        ethernet-ring ring-2 {
            east-interface {
                control-channel {
                    ge-3/2/4.0;
                    vlan 101;
                }
            }
            west-interface {
                control-channel {
                    ge-5/2/3.0;
                    vlan 101;
                }
            }
            data-channel {
                vlan [ 500 600 ];
            }
        }
    }
}
```

```

    }
  }
  routing-instances {
    vs {
      instance-type virtual-switch;
      interface ge-3/2/4.0;
      interface ge-5/2/3.0;
      bridge-domains {
        bd100 {
          vlan-id 100;
        }
        bd101 {
          vlan-id 101;
        }
        bd200 {
          vlan-id 200;
        }
        bd300 {
          vlan-id 300;
        }
        bd500 {
          vlan-id 500;
        }
        bd600 {
          vlan-id 600;
        }
      }
    }
  }
}

```

Configuring ERP on CS2

CLI Quick Configuration

To quickly configure CS2 for ERP, copy the following commands and paste them into the switch terminal window of CS2:

```

[edit]
set interfaces ge-2/0/4 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/4 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-2/0/8 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/8 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-2/0/4.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-2/0/8.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
set protocols protection-group ethernet-ring ring-2 ring-protection-link-owner
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-2/0/4.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-2/0/8.0
set protocols protection-group ethernet-ring ring-2 west-interface ring-protection-link-end
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set bridge-domains bd100 vlan-id 100
set bridge-domains bd101 vlan-id 101
set bridge-domains bd200 vlan-id 200
set bridge-domains bd300 vlan-id 300

```

```
set bridge-domains bd500 vlan-id 500
set bridge-domains bd600 vlan-id 600
```

Step-by-Step Procedure

To configure ERP on CS2:

1. Configure the trunk interface **ge-2/0/4** to connect CS2 to CS1 and trunk interface **ge-2/0/8** to connect CS2 to CS1, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```
[edit interfaces]
user@cs2# set ge-2/0/4 unit 0 family bridge interface-mode trunk
user@cs2# set ge-2/0/4 unit 0 family bridge vlan-id-list 100-1000
user@cs2# set ge-2/0/8 unit 0 family bridge interface-mode trunk
user@cs2# set ge-2/0/8 unit 0 family bridge vlan-id-list 100-1000
```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**, and configure **ring-2** as the ring protection link owner:



NOTE: Always configure the east-interface statement first, before configuring the west-interface statement.

```
[edit protection-group]
user@cs2# set ethernet-ring ring-1 east-interface control-channel ge-2/0/4.0
user@cs2# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@cs2# set ethernet-ring ring-1 west-interface control-channel ge-2/0/8.0
user@cs2# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@cs2# set ethernet-ring ring-2 data-channel vlan [200,300]
user@cs2# set ethernet-ring ring-2 east-interface control-channel ge-2/0/4.0
user@cs2# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@cs2# set ethernet-ring ring-2 ring-protection-link-owner
user@cs2# set ethernet-ring ring-2 west-interface control-channel ge-2/0/8.0
user@cs2# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@cs2# set ethernet-ring ring-2 west-interface ring-protection-link-end
user@cs2# set ethernet-ring ring-2 data-channel vlan [500,600]
```

3. Configure the routing instance, the bridge domains, and the VLAN IDs associated with each bridge domain:

```
[edit bridge-domains]
user@cs2# set bd100 vlan-id 100
user@cs2# set bd101 vlan-id 101
user@cs2# set bd200 vlan-id 200
user@cs2# set bd300 vlan-id 300
user@cs2# set bd500 vlan-id 500
user@cs2# set bd600 vlan-id 600
```

Results Check the results of the configuration:

```
user@cs2> show configuration
interfaces {
  ge-2/0/4 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
      }
    }
  }
}
```

```
}
ge-2/0/8 {
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 100-1000;
    }
  }
}
protocols {
  protection-group {
    ethernet-ring ring-1 {
      east-interface {
        control-channel {
          ge-2/0/4.0;
          vlan 100;
        }
      }
      west-interface {
        control-channel {
          ge-2/0/8.0;
          vlan 100;
        }
      }
      data-channel {
        vlan [200, 300];
      }
    }
  }
}
ethernet-ring ring-2 {
  east-interface {
    control-channel {
      ge-2/0/4.0;
      vlan 101;
    }
  }
  west-interface {
    control-channel {
      ge-2/0/8.0;
      vlan 101;
    }
  }
  ring-protection-link-end;
}
data-channel {
  vlan [500, 500];
}
}
bridge-domains {
  bd100 {
    vlan-id 100;
  }
  bd101 {
    vlan-id 101;
  }
  bd200 {
```

```

        vlan-id 200;
    }
    bd300 {
        vlan-id 300;
    }
    bd500 {
        vlan-id 500;
    }
    bd600 {
        vlan-id 600;
    }
}
}

```

Configuring ERP on AS1

CLI Quick Configuration To quickly configure AS1 for ERP, copy the following commands and paste them into the switch terminal window of AS1:

```

[edit]
set interfaces ge-2/0/5 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/5 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-2/1/1 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/1 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-2/0/5.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-2/1/1.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-2/0/5.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-2/1/1.0
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set bridge-domains bd100 vlan-id 100
set bridge-domains bd101 vlan-id 101
set bridge-domains bd200 vlan-id 200
set bridge-domains bd300 vlan-id 300
set bridge-domains bd500 vlan-id 500
set bridge-domains bd600 vlan-id 600

```

Step-by-Step Procedure To configure ERP on AS1:

1. Configure the trunk interface **ge-2/0/5** to connect CS2 to CS1 and trunk interface **ge-2/1/1** to connect CS2 to CS1, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```

[edit interfaces]
user@as1# set ge-2/0/5 unit 0 family bridge interface-mode trunk
user@as1# set ge-2/0/5 unit 0 family bridge vlan-id-list 100-1000
user@as1# set ge-2/1/1 unit 0 family bridge interface-mode trunk
user@as1# set ge-2/1/1 unit 0 family bridge vlan-id-list 100

```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**:



NOTE: Always configure the east-interface statement first, before configuring the west-interface statement.

```

[edit protection-group]
user@as1# set ethernet-ring ring-1 east-interface control-channel ge-2/0/5.0
user@as1# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@as1# set ethernet-ring ring-1 west-interface control-channel ge-2/1/1.0
user@as1# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@as1# set ethernet-ring ring-2 east-interface control-channel ge-2/0/5.
user@as1# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@as1# set ethernet-ring ring-2 west-interface control-channel ge-2/1/1.0
user@as1# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@as1# set ethernet-ring ring-2 data-channel vlan [500, 600]
3. Configure the routing instance, the bridge domains, and the VLAN IDs associated
   with each bridge domain:

[edit bridge-domains]
user@as1# set bd100 vlan-id 100
user@as1# set bd101 vlan-id 101
user@as1# set bd200 vlan-id 200
user@as1# set bd300 vlan-id 300
user@as1# set bd500 vlan-id 500
user@as1# set bd600 vlan-id 600

```

Results Check the results of the configuration:

```

user@as1> show configuration
interfaces {
  ge-2/0/5 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000
      }
    }
  }
  ge-2/1/1 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000
      }
    }
  }
}
protocols {
  protection-group {
    ethernet-ring ring-1 {
      east-interface {
        control-channel {
          ge-2/0/5.0;
          vlan 100;
        }
      }
      west-interface {
        control-channel {
          ge-2/1/1.0;
          vlan 100;
        }
      }
      data-channel {

```

```
        vlan [200, 300];
    }
}
}
protection-group {
  ethernet-ring ring-2 {
    east-interface {
      control-channel {
        ge-2/0/5.0;
        vlan 101;
      }
    }
    west-interface {
      control-channel {
        ge-2/1/1.0;
        vlan 101;
      }
    }
    data-channel {
      vlan [500, 600];
    }
  }
}
bridge-domains {
  bd100 {
    vlan-id 100;
  }
  bd101 {
    vlan-id 101;
  }
  bd200 {
    vlan-id 200;
  }
  bd300 {
    vlan-id 300;
  }
  bd500 {
    vlan-id 500;
  }
  bd600 {
    vlan-id 600;
  }
}
}
```

Verification

To confirm that the ERP configuration for multiple ring instances is operating, perform these tasks:

- [Verifying the Ethernet Protection Ring on CS1 on page 81](#)
- [Verifying the Data Channel CS1 on page 81](#)
- [Verifying the VLANs on CS1 on page 82](#)

- [Verifying the Ethernet Protection Ring on CS2 on page 82](#)
- [Verifying the Data Channel CS2 on page 83](#)
- [Verifying the VLANs on CS2 on page 83](#)
- [Verifying the Ethernet Protection Ring on AS1 on page 84](#)
- [Verifying the Data Channels on AS1 on page 84](#)
- [Verifying the VLANs on AS1 on page 85](#)

Verifying the Ethernet Protection Ring on CS1

Purpose Verify that ERP is enabled on CS1.

Action Show the status of the ring automatic protection switching (R-APS) messages to determine if there is a ring failure:

```
user@cs1>show protection-group ethernet-ring aps
```

Ethernet Node ID	Ring Name	Request/state	No Flush	Ring Protection	Link Blocked	Originator	Remote
ring-1		NR	No	Yes	Yes	Yes	
ring-2		NR	No	Yes	Yes	No	
13:22:af:31:fc:00							

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** displays **yes**, indicating that this node is an R-APS originator.

Verifying the Data Channel CS1

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the Spanning Tree Protocol (STP) index number):

```
user@cs1>show protection-group ethernet-ring data-channel
```

```
Ethernet ring data channel parameters for protection group ring-1
```

Interface	STP index	Forward State
ge-3/2/4	122	forwarding
ge-5/2/3	123	forwarding

```
Ethernet ring data channel parameters for protection group ring-2
```

Interface	STP index	Forward State
-----------	-----------	---------------

ge-3/2/4	124	discarding
ge-5/2/3	125	forwarding

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward State** shows whether the data channel is **forwarding** or **discarding** traffic.

Verifying the VLANs on CS1

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@cs1>show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group ring-1
```

Interface	Vlan	STP Index	Bridge Domain
ge-3/2/4	200	122	vs/bd200
ge-5/2/3	200	123	vs/bd200
ge-3/2/4	300	122	vs/bd300
ge-5/2/3	300	123	vs/bd300

```
Ethernet ring IFBD parameters for protection group ring-2
```

Interface	Vlan	STP Index	Bridge Domain
ge-3/2/4	500	124	vs/bd500
ge-5/2/3	500	125	vs/bd500
ge-3/2/4	600	124	vs/bd600
ge-5/2/3	600	125	vs/bd600

Meaning The output displayed shows the ring interfaces **ge-3/2/4** and **ge-5/2/3** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 122** and **123** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 124** and **125** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load balancing.

Verifying the Ethernet Protection Ring on CS2

Purpose Verify that ERP is enabled on CS2.

Action Show the status of the ring APS (R-APS) messages to determine if there is a ring failure:

```
user@cs2>show protection-group ethernet-ring aps
```

Ethernet Ring Name Node ID	Request/state	No Flush	Ring Protection	Originator	Remote
Ring-1 00:21:59:03:ff:d0	NR	No	No	No	
Ring-2	NR	No	Yes	Yes	

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** displays **yes**, indicating that this node is an R-APS originator. The **Originator** field for **ring-2** displays **No**, indicating that this node is not an R-APS originator.

Verifying the Data Channel CS2

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the STP index number):

```
user@cs2> show protection-group ethernet-ring data-channel
Ethernet ring data channel parameters for protection group ring-1
```

Interface	STP index	Forward State
ge-2/0/4	44	forwarding
ge-2/0/8	45	forwarding

```
Ethernet ring data channel parameters for protection group ring-2
```

Interface	STP index	Forward State
ge-2/0/4	46	forwarding
ge-2/0/8	47	discarding

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward State** shows whether the data channel is **forwarding** or **discarding** traffic.

Verifying the VLANs on CS2

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@cs2> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group ring-1
```

Interface	Vlan	STP Index	Bridge Domain
ge-2/0/4	200	44	default-switch/bd200
ge-2/0/8	200	45	default-switch/bd200
ge-2/0/4	300	44	default-switch/bd300
ge-2/0/8	300	45	default-switch/bd300

```
Ethernet ring IFBD parameters for protection group ring-2
```

Interface	Vlan	STP Index	Bridge Domain
ge-2/0/4	500	46	default-switch/bd500
ge-2/0/8	500	47	default-switch/bd500
ge-2/0/4	600	46	default-switch/bd600
ge-2/0/8	600	47	default-switch/bd600

Meaning The output displayed shows the ring interfaces **ge-2/0/4** and **ge-2/0/8** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 44** and **45** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 46** and **47** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load balancing.

Verifying the Ethernet Protection Ring on AS1

Purpose Verify that ERP is enabled on AS1.

Action Show the status of the ring APS (R-APS) messages to determine if there is a ring failure:

```
user@as1> show protection-group ethernet-ring aps
```

Ethernet Ring Name Node ID	Request/state	No Flush	Ring Protection	Originator	Remote
Ring-1 00:21:59:03:ff:d0	NR	No	Link Blocked Yes	No	
Ring-2 13:22:af:31:fc:00	NR	No	Yes	No	

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** and **ring-2** displays **No**, indicating that this node is not the R-APS originator.

Verifying the Data Channels on AS1

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the STP index number):

```
user@as1> show protection-group ethernet-ring data-channel
Ethernet ring data channel parameters for protection group ring-1
```

Interface	STP index	Forward State
ge-2/0/5	22	forwarding
ge-2/1/1	23	forwarding

```
Ethernet ring data channel parameters for protection group ring-2
```

Interface	STP index	Forward State
ge-2/0/5	24	forwarding
ge-2/1/1	25	forwarding

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward State** shows whether the data channel is **forwarding** or **discarding** traffic. All data channels are forwarding traffic.

Verifying the VLANs on AS1

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@as1>show protection-group ethernet-ring vlan
```

Ethernet ring IFBD parameters for protection group ring-1

Interface	Vlan	STP Index	Bridge Domain
ge-2/0/5	200	22	default-switch/bd200
ge-2/1/1	200	23	default-switch/bd200
ge-2/0/5	300	22	default-switch/bd300
ge-2/1/1	300	23	default-switch/bd300

Ethernet ring IFBD parameters for protection group ring-2

Interface	Vlan	STP Index	Bridge Domain
ge-2/0/5	500	24	default-switch/bd500
ge-2/1/1	500	25	default-switch/bd500
ge-2/0/5	600	24	default-switch/bd600
ge-2/1/1	600	25	default-switch/bd600

Meaning The output displayed shows the ring interfaces **ge-2/0/5** and **ge-2/1/1** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 22** and **23** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 24** and **25** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load-balancing.

Related Documentation

- [Ethernet Ring Protection Using Ring Instances for Load Balancing on page 15](#)
- [Ethernet Ring Protection on page 13](#)

Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation

Under normal operating conditions, when Ethernet ring protection is configured correctly, the ring protection link (RPL) owner (Router 1 in the configuration example) will see the following:

Router 1 Operational Commands (Normal Ring Operation)

```
user@router1> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg101              NR           No      Yes
```

```
Originator Remote Node ID
Yes
```

Note that the ring protection link is blocked and the node is marked as the originator of the protection.

```
user@router1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	discarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

```
Signal Failure Admin State
Clear          IFF ready
Clear          IFF ready
```

Note that the protection interface is discarding while the other interface is forwarding.

```
user@router1> show protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg101         idle     NR-RB Yes
```

```
Restore Timer Guard Timer Operation state
disabled      disabled    operational
```

Note that Router 1 is the owner and timers are disabled.

```
user@router1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent           : 1
RAPS received       : 0
Local SF happened:   : 0
Remote SF happened:   : 0
NR event happened:    : 0
NR-RB event happened: : 1
```

Note that only minimal RAPS messages have been sent to establish the ring.

Under normal operating conditions, the other routers on the ring (Router 2 and Router 3) will see the following similar output:

Router 2 and Router 3 Operational Commands (Normal Ring Operation)

```
user@router2> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg102              NR           No      Yes
```

```
Originator Remote Node ID
No          00:01:01:00:00:01
```

Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel Forward State Ring Protection Link End
ge-1/2/1     ge-1/2/1.1      forwarding   No
ge-1/0/2     ge-1/0/2.1      forwarding   No

Signal Failure Admin State
Clear         IFF ready
Clear         IFF ready
```

Note that both interfaces are forwarding. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg102         idle      NR-RB      No

Restore Timer Guard Timer Operation state
disabled      disabled  operational
```

Note that Router 2 is not the owner. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent : 0
RAPS received : 1
Local SF happened: : 0
Remote SF happened: : 0
NR event happened: : 0
NR-RB event happened: : 1
```

Router 3 will see almost identical information.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Ring Protection on page 13](#)
- [Example: Configuring Ethernet Ring Protection for MX Series Routers on page 61](#)
- [Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 87](#)

Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition

This section assumes that Ethernet ring protection is configuring correctly, that Router 1 is the ring protection link (RPL) owner, and that there is a link failure between Router 2 and Router 3 in the configuration example.

Router 1 Operational Commands (Ring Failure Condition)

```
user@router1> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg101             SF             NO      No
```

```
Originator Remote Node ID
No          00:01:02:00:00:01
```

Note that the ring protection link is no longer blocked and the node is no longer marked as originator.

```
user@router1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	forwarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

```
Signal Failure Admin State
Clear          IFF ready
Clear          IFF ready
```

Note that the protection interface is now forwarding (so is the other interface).

```
user@router1> show protection-group ethernet-ring node-state
how protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg101          protected SF Yes
```

```
Restore Timer Quard Timer Operation state
disabled      disabled operational
```

Note that Router 1 has recorded the span failure (SF).

```
user@router1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent : 1
RAPS received : 1
Local SF happened: : 0
Remote SF happened: : 1
NR event happened: : 0
NR-RB event happened: : 1
```

Note that the R-APS messages have recorded the remote failure.

Under a failure condition, the other routers on the ring (Router 2 and Router 3) will see the following similar output:

Router 2 and Router 3 Operational Commands (Failure Condition)

```
user@router2> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg102             SF           No      No
```

```
Originator Remote Node ID
Yes         00:00:00:00:00:00
```

Note the failure event (SF). Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/2/1	ge-1/2/1.1	forwarding	No
ge-1/0/2	ge-1/0/2.1	discarding	No

```
Signal Failure Admin State
Clear          IFF ready
set            IFF ready
```

Note that the failed interface (**ge-1/0/2.1**) is not forwarding. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring node-state
Ethernet ring  APS State  Event          Ring Protection Link Owner
pg102          idle      NR-RB         No
```

```
Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

Note that Router 2 is not the owner. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                : 1
RAPS received            : 1
Local SF happened:       : 1
Remote SF happened:      : 0
NR event happened:       : 0
NR-RB event happened:    : 1
```

Note that the R-APS messages have recorded the remote failure. Router 3 will see almost identical information.

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Ring Protection on page 13](#)
- [Example: Configuring Ethernet Ring Protection for MX Series Routers on page 61](#)
- [Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 85](#)

PART 3

Administration

- [Managing Ethernet Frame Delay Measurements on page 93](#)

CHAPTER 10

Managing Ethernet Frame Delay Measurements

- [Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements on page 93](#)
- [Triggering an Ethernet Frame Delay Measurements Session on page 94](#)

Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging Service Level Agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

Before you can perform Ethernet frame delay measurements on MX Series routers, you must have done the following:

- Configured Ethernet OAM and CFM correctly
- Prepared the measurement between two compatibly configured MX Series routers
- Enabled the distributed periodic packet management daemon (ppmd)
- Avoided trying to perform Ethernet frame delay measurement on aggregated Ethernet or pseudowire interfaces, which are not supported
- Made sure the hardware-assisted timestamping is supported if that feature is configured

At the end of this configuration, you create two MX Series routers that can perform and display Ethernet frame delay measurements on Ethernet interfaces using optional hardware timestamping. By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

To configure hardware-assisted timestamping:

1. To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        performance-monitoring {
          hardware-assisted-timestamping; # Enable timestamping in hardware.
        }
      }
    }
  }
}
```

2. Ethernet frame delay measurement requires that distributed PPMD is enabled. Before you can gather statistics for Ethernet frame delay measurement, you must make sure that PPMD is configured properly. Without distributed PPMD, delay measurement results are not valid.

To perform Ethernet frame delay measurement, make sure that the following configuration statement is *NOT* present:

```
[edit routing-options]
ppm {
  no-delegate-processing; # This turns distributed PPMD OFF.
}
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Triggering an ETH-DM Session on page 20](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

Triggering an Ethernet Frame Delay Measurements Session

Before Ethernet frame delay measurement statistics can be displayed, they must be collected. To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement (one-way | two-way) (remote-mac-address | mep identifier) maintenance-domain name maintenance-association ma-id [count count] [wait time]** operational command.

The fields for this command are described in [Table 3 on page 21](#).

Table 6: Monitor Ethernet Delay Command Parameters

Parameter	Parameter Range	Description
one-way or two-way	NA	Perform a one-way or two-way (round-trip) delay measurement.
remote-mac-address	Unicast MAC address	Send delay measurement frames to the destination unicast MAC address (use the format xx:xx:xx:xx:xx:xx). Multicast MAC addresses are not supported.
mep identifier	1–8191	The MEP identifier to use for the measurement. The discovered MAC address for this MEP identifier is used.
maintenance-domain name	Existing MD name	Specifies an existing maintenance domain (MD) to use for the measurement.
maintenance-association ma-id	Existing MA identifier	Specifies an existing maintenance association (MA) identifier to use for the measurement.
count count	1–65535 (default: 10)	(Optional) Specifies the number of Ethernet frame delay frames to send. The default is 10.
wait time	1–255 seconds (default: 1)	(Optional) Specifies the number of seconds to wait between frames. The default is 1 second.

If you attempt to monitor delays to a nonexistent MAC address, you must exit the application manually using ^C:

```
user@host> monitor ethernet delay-measurement two-way 00:11:22:33:44:55
Two-way ETH-DM request to 00:11:22:33:44:55, Interface ge-5/2/9.0
^C
--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 0
Average delay: 0 usec, Average delay variation: 0 usec
Best case delay: 0 usec, Worst case delay: 0 usec
```

Related Documentation

- [Ethernet OAM Feature Guide for MX Series Routers](#)
- [Ethernet Frame Delay Measurements on page 7](#)
- [Configuring MEP Interfaces to Support ETH-DM on page 19](#)
- [Viewing ETH-DM Statistics on page 21](#)
- [Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 41](#)
- [Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 46](#)
- [Configuring ETH-DM with Untagged Interfaces on page 50](#)

