



MPLS Management Suite

1.2

USER'S GUIDE

P/N 300-002-533

REV A02

OL-8971-01

EMC Corporation

Corporate Headquarters:

Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright © 2003-2005 by EMC Corporation ("EMC"). All rights reserved.

EMC Corporation believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OR MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL EMC CORPORATION BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF BUSINESS, LOSS OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND, ARISING FROM ANY ERROR IN THIS PUBLICATION.

The EMC Smarts software products are covered by one or more of U.S. Patent Nos. or pending patent applications assigned to EMC Corporation.

"EMC," "InCharge," the InCharge logo, "SMARTS," the SMARTS logo, "Graphical Visualization," "Authentic Problem," "Codebook Correlation Technology," "Instant Results Technology," "InCharge Viewlet," and "Dashboard Viewlet" are trademarks or registered trademarks of EMC Corporation. All other brand or product names are trademarks or registered trademarks of their respective companies or organizations.

Additional copyright notices and license terms applicable to the software product are set forth in the Third-Party Copyright Read Me file included on the accompanying software media.

Contents

Preface	ix
Intended Audience	ix
Prerequisites	ix
Document Organization	x
Documentation Conventions	xi
MPLS Management Suite Installation Directory	xii
MPLS Management Suite Products	xii
Additional Resources	xii
Command Line Programs	xii
Documentation	xiii
Technical Support	xiii
EMC Powerlink	xiv
1 Introduction	1
Discovery	2
Availability Manager Discovery	2
MPLS Manager Discovery	3
Remote Ping Functionality	7
Monitoring	8
Analysis	8
Notifications	8
2 MPLS and VPN Elements and Their Failures	9
MPLS and VPN Elements and Relationships	10
Summary of Events for L3VPN Networks	11
Summary of Events for L2VPN Networks	11
MPLS Core Elements	13
MPLSService	13

LSP	14
LSPHop	15
LSPInSegment	17
LSPOutSegment	17
L3VPN Elements	19
VPN (Layer 3)	20
VRF	21
RouteTarget	24
L2VPN Elements	25
VPN (Layer 2)	25
Forwarder	26
PseudoWire	28
LdpProtocolEndpoint	30
LdpAdjacency	31
Provisioning System Adapter Events	33
Underlying Transport Network Failures	33
3 MPLS Cross-Domain Impact Correlation Analysis	35
Impact Analysis Overview	35
Impact Analysis Models	36
Impact Analysis Events	40
L3VPN Domain Impact Events	40
L2VPN Domain Impact Events	40
MPLS Domain Impact Events	41
Impact Analysis Examples for the L3VPN Domain	42
L3VPN Impact Example 1: CE-PE Router NetworkConnection Down	42
L3VPN Impact Example 2: PE Router Down	44
Impact Analysis Examples for the L2VPN Domain	46
4 Viewing MPLS Notifications, Maps, and Containment	49
Viewing MPLS Notifications	50
Opening an MPLS Notification Properties Dialog Box	50
MPLS Notification Properties	51
Viewing MPLS Topology in Maps	52

Opening an MPLS Topology Map	53
MPLS Topology Map Graphical Representations	53
MPLS Map Types	55
Viewing MPLS Containment	65
Opening an MPLS Containment Dialog	66
MPLS Containment Tab Pages	67
5 Customizing Groups and Settings	69
Default Groups and Settings	70
Forwarders	72
LdpProtocolEndpoint	73
VRFs	74
CE to CE Pings	76
PE to CE Ping Setting	78
PE to PE Ping Setting	80
PE to Unmanaged CE Ping Settings	82
PE to VRF Ping Setting	84
System Write Community Strings	86
Default Threshold Groups and Settings	87
Opening the Polling and Thresholds Console	87
Layout of the Polling and Thresholds Console	88
Polling and Thresholds Console Toolbar Buttons	89
Working With Groups and Settings	90
How Managed Elements Are Assigned to Groups	90
Modifying the Properties of a Group	90
Adding or Removing Group Settings	91
Modifying the Priority of Groups	91
Editing Matching Criteria	92
Modifying the Parameters of a Setting	93
Creating New Groups	94
6 Remote Ping Functionality	97
Types of Remote Ping Requests	97
How Remote Pings are Generated	98

Examples of Remote Pings	99
Remote Ping Example 1: PE to PE	99
Remote Ping Example 2: CE to CE	99
Remote Ping Example 3: PE to Remote CE	100
Remote Ping Example 4: PE to Local CE	100
Remote Ping Example 5: PE to Unmanaged CE	101
Remote Ping Example 6: PE to VRF	101
More about VRF Ping	102
Remote Ping Elements	102
Attributes for RemotePing Elements	103
Relationships for RemotePing Elements	104
Remote Ping Impact Analysis	104
Viewing Periodic RemotePing Information	105
Notification Log	105
Notification Properties Window	106
Topology Browser	107
Issuing an On-Demand Remote Ping	108
Using the Set Ping Source Server Tool	108
Using the Who's My Ping Source Server Tool	109
Using the Remote Ping Server Tool	110
Using the VRF Ping Server Tool	110
Using the Repeat Remote Ping Server Tool	111
Log Files	112
A MPLS Terminology	113
B Root-Cause Notifications from Availability Manager	125
C Polling for Analysis	127
SNMP Poller	127
Just-In-Time Polling	128
Request-Consolidation Polling	128

D Wildcard Patterns **129**

Index **133**

Preface

This document provides detailed information about the EMC Smarts MPLS Manager. The MPLS Manager, in conjunction with the EMC Smarts IP Availability Manager (Availability Manager), diagnoses connectivity failures in Multiprotocol Label Switching (MPLS) networks and MPLS Layer 3 and Layer 2 Virtual Private Networks (VPNs), and then sends the results of its analysis to the EMC Smarts Global Manager (Service Assurance).

The MPLS Manager also supports remote ping functionality, allowing for periodic or on-demand pings from various Layer 3 VPN elements to other Layer 3 VPN elements.

Intended Audience

This document is intended to be read by operators receiving and acting upon MPLS Manager notifications, by system administrators configuring and using the MPLS Manager, and by IT managers seeking to better understand the value of the MPLS Manager.

Prerequisites

Before you perform the procedures in this document, the following EMC Smarts software must be installed:

- Availability Manager
- Service Assurance Manager (Global Manager)
- Global Console
- MPLS Manager

For information about installing these products, see the *EMC Smarts IP Management Suite Installation Guide*, the *EMC Smarts Service Assurance Management Suite Installation Guide*, and the *EMC Smarts MPLS Management Suite Installation Guide*.

Document Organization

This document consists of the following chapters and appendices.

Table 1: Document Organization

1. INTRODUCTION	Describes the concepts of managing MPLS network connectivity using the MPLS Manager.
2. MPLS AND VPN ELEMENTS AND THEIR FAILURES	Describes the MPLS and VPN elements managed by the MPLS Manager and identifies the root-cause problems and symptomatic events for each element type.
3. MPLS CROSS-DOMAIN IMPACT CORRELATION ANALYSIS	Describes how the MPLS Manager correlates failures in the MPLS and VPN domains with failures in the transport network domain to perform MPLS impact analysis.
4. VIEWING MPLS NOTIFICATIONS, MAPS, AND CONTAINMENT	Describes how to use the Global Console to view MPLS notifications, topology maps, and containment information for the MPLS Manager.
5. CUSTOMIZING GROUPS AND SETTINGS	Provides procedures for configuring SNMP polling and for configuring periodic remote ping instances.
6. REMOTE PING FUNCTIONALITY	Describes the remote ping functionality and provides procedures for configuring and using remote ping.
A. MPLS TERMINOLOGY	Describes basic terms and concepts for MPLS and VPN networks.
B. ROOT-CAUSE NOTIFICATIONS FROM AVAILABILITY MANAGER	Lists the notifications that the MPLS Manager receives from the Availability Manager.
C. POLLING FOR ANALYSIS	Describes the SNMP polling engine used by the MPLS Manager for correlation analysis.
D. WILDCARD PATTERNS	Describes the wildcards used to create matching patterns.

Documentation Conventions

Several conventions may be used in this document as shown in Table 2.

Table 2: Documentation Conventions

CONVENTION	EXPLANATION
<code>sample code</code>	Indicates code fragments and examples in Courier font
keyword	Indicates commands, keywords, literals, and operators in bold
%	Indicates C shell prompt
#	Indicates C shell superuser prompt
<parameter>	Indicates a user-supplied value or a list of non-terminal items in angle brackets
[option]	Indicates optional terms in brackets
<i>/InCharge</i>	Indicates directory path names in italics
<i>yourDomain</i>	Indicates a user-specific or user-supplied value in bold, italics
<i>File > Open</i>	Indicates a menu path in italics
▼▲	Indicates a command is wrapped over one or more lines. The command must be typed as one line.

Directory path names are shown with forward slashes (/). Users of the Windows operating systems should substitute back slashes (\) for forward slashes.

Also, if there are figures illustrating consoles in this document, they represent the consoles as they appear in Windows. Under UNIX, the consoles appear with slight differences. For example, in views that display items in a tree hierarchy such as the Topology Browser, a plus sign displays for Windows and an open circle displays for UNIX.

Finally, unless otherwise specified, the term InCharge Manager is used to refer to EMC Smarts programs such as Domain Managers, Global Managers, and adapters.

MPLS Management Suite Installation Directory

In this document, the term **BASEDIR** represents the location where EMC Smarts software is installed.

- For UNIX, this location is: `/opt/InCharge<n>/<productsuite>`.
- For Windows, this location is: `C:\InCharge<n>\<productsuite>`.

The `<n>` represents the EMC Smarts software platform version number. The `<productsuite>` represents the InCharge product suite to which the product belongs. For example, on UNIX operating systems, MPLS Manager is installed to `/opt/InCharge6/MPLS/smarts` by default. On Windows operating systems, this product is installed to `C:\InCharge6\MPLS\smarts` by default. This location is referred to as **BASEDIR**/`smarts`.

Optionally, you can specify the root of **BASEDIR** to be something other than `/opt/InCharge6` (on UNIX) or `C:\InCharge6` (on Windows), but you cannot change the `<productsuite>` location under the root directory.

For more information about the directory structure of EMC Smarts software, refer to the *EMC Smarts System Administration Guide*.

MPLS Management Suite Products

The MPLS Management Suite offers the following products:

- MPLS Manager
- EMC Smarts Adapter for Cisco ISC
- Perl API

Additional Resources

In addition to this document, EMC Corporation provides the following resources.

Command Line Programs

Descriptions of command line programs are available as HTML pages. The `index.html` file, which provides an index to the various commands, is located in the **BASEDIR**/`smarts/doc/html/usage` directory.

Documentation

Readers of this document may find other documentation (also available in the **BASEDIR**/*smarts/doc/pdf* directory) helpful.

EMC Smarts Documentation

The following documents are product independent and thus relevant to users of all EMC Smarts products:

- *EMC Smarts Documentation Roadmap*
- *EMC Smarts System Administration Guide*
- *EMC Smarts ASL Reference Guide*
- *EMC Smarts Perl Reference Guide*

MPLS Management Suite Documentation

The following documents are relevant to users of the MPLS Management Suite product suite:

- *EMC Smarts MPLS Management Suite Installation Guide*
- *EMC Smarts MPLS Manager User's Guide*
- *EMC Smarts MPLS Manager Configuration Guide*
- *EMC Smarts MPLS Manager Discovery Guide Supplement*
- *EMC Smarts Adapter for Cisco ISC User's Guide*
- *EMC Smarts MPLS Management Suite Release Notes*

Refer to the *EMC Smarts Documentation Roadmap* for documentation resources provided with other EMC Smarts product suites.

Technical Support

For questions about technical support, call your local sales office or service provider. For service, call one of the following numbers:

United States: 800.782.4362 (SVC.4EMC)

Canada: 800.543.4782 (543.4SVC)

Worldwide: 508.497.7901

EMC Powerlink

EMC Powerlink is the EMC Corporation's secure extranet for customers and partners. Powerlink is an essential tool for obtaining web-based support from the EMC Corporation. Powerlink can be used to submit service or information requests (tickets) and monitor their progress, to review the knowledgebase for known problems and solutions, and to download patches and SmartPacks.

From training on EMC products and technologies, to online support, product announcements, software registration, technical white papers, interoperability information, and a range of configuration tools, Powerlink offers resources unavailable elsewhere.

For quickest access when you do not already have a Powerlink account, ask your EMC representative for the access code for your company and register at the Powerlink site. Visit the EMC Powerlink website at:

<http://powerlink.emc.com>

Introduction

Multiprotocol Label Switching (MPLS) provides IP networks with the kind of traffic management and connection-oriented quality of service found in networks like Asynchronous Transfer Mode (ATM) and Frame Relay. MPLS enhances network performance by introducing virtual circuits called Label Switched Paths (LSPs) to IP networks: Packets are *switched* rather than *routed* through the network. And because the fundamental principles of virtual circuits are based on traffic separation and segmentation, MPLS is ideal for building provider-provisioned Layer 3 (L3) and Layer 2 (L2) Virtual Private Networks (VPNs).

The MPLS Manager, working with the Availability Manager and other components of the MPLS Manager architecture, performs the following major functions:

- For MPLS-based L3VPN networks, discovers and monitors network, MPLS, and L3VPN elements
- For MPLS-based L2VPN networks, discovers and monitors network, MPLS, and L2VPN elements
- Correlates underlying network problems with MPLS, L3VPN, and L2VPN impairments
- Identifies configuration and other errors that occur when deploying and maintaining MPLS and VPN networks
- Performs root-cause analysis and reports its analysis results to the Global Manager

In addition, the MPLS Manager provides remote ping functionality, which allows for periodic and on-demand pings from various L3VPN elements to other L3VPN elements.

The MPLS Manager architecture is illustrated and described in the *EMC Smarts MPLS Manager Configuration Guide*.

Discovery

The MPLS Manager works with the Availability Manager to discover the logical and physical elements in the transport domain, the MPLS domain, and the VPN domain.

Availability Manager Discovery

In the transport domain, the Availability Manager discovers the Layer 2 and Layer 3 network element connectivity between Provider (P) routers, Provider Edge (PE) routers, and Customer Edge (CE) routers (Figure 1). In addition, it discovers and models the Virtual Local Area Networks (VLANs).

The Availability Manager uses the discovered topology to model the network, and uses SNMP polling and traps to diagnose and pinpoint the root cause of network failures. The Availability Manager sends the analysis results along with topology and event information to the Global Manager, and sends router topology and event information to the MPLS Manager.

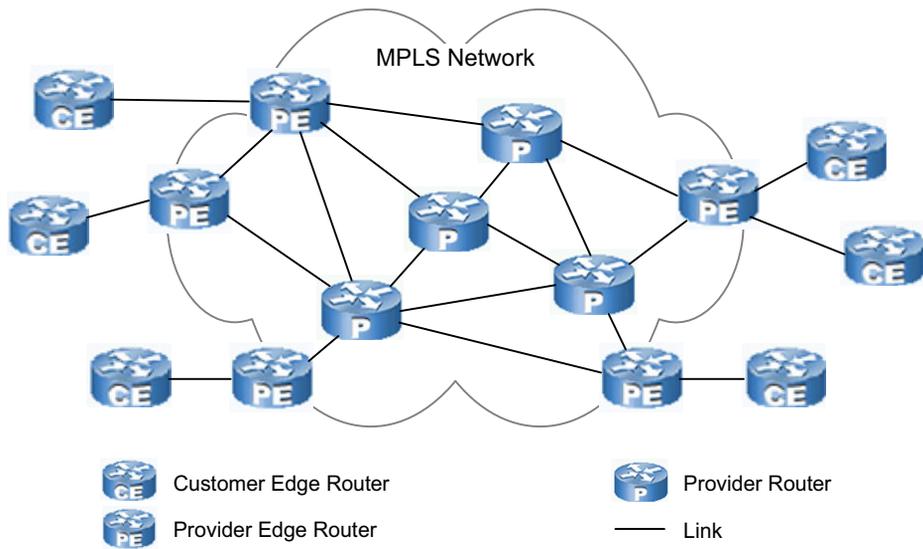


Figure 1: The Transport Domain Discovered by the Availability Manager

MPLS Manager Discovery

The MPLS Manager discovers the MPLS and VPN logical topology and models that topology in its repository.

MPLS Domain

For the MPLS domain, shown in Figure 2, the MPLS Manager maps the MPLS and VPN topology to the router topology discovered by the Availability Manager.

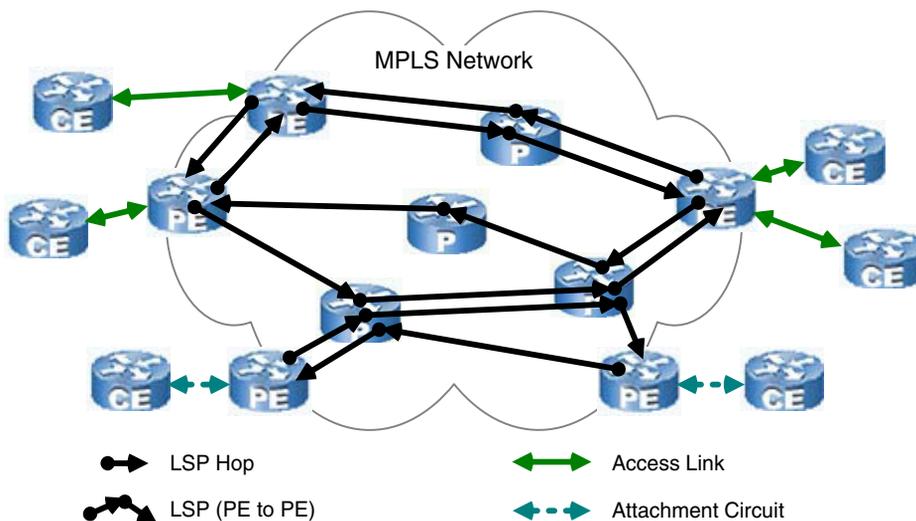


Figure 2: The MPLS Domain Discovered by the MPLS Manager

L3VPN Domain

For the L3VPN domain, shown in Figure 3, the MPLS Manager discovers the VPN Routing and Forwarding (VRF) instances that manage the VPN routes, and the VPNs themselves.

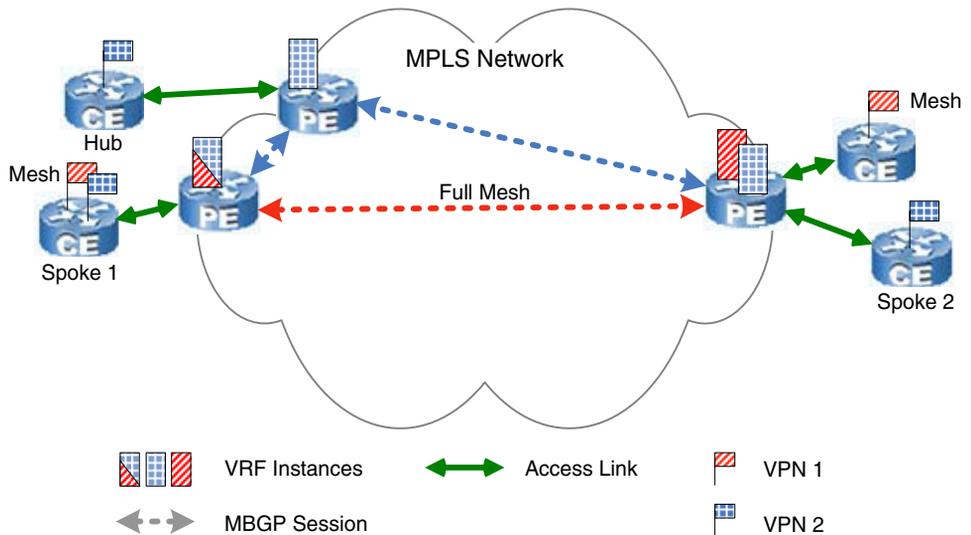


Figure 3: The L3VPN Domain Discovered by the MPLS Manager

L2VPN Domain

The MPLS Manager discovers elements in L2VPNs domains that are based on either the Martini or Kompella methodologies. Martini-implemented L2VPNs use Label Distribution Protocol (LDP) as their signaling protocol, and Kompella-implemented L2VPNs use Multiprotocol Border Gateway Protocol (MBGP) as their signaling protocol.

For the L2VPN domain, the MPLS Manager supports the Virtual Private Wire Service (VPWS), which is a point-to-point L2VPN service that is formed by combining the following:

- Source Attachment Circuit—The circuit or virtual circuit that attaches a CE device to a PE device; for example, a VLAN or an Ethernet port.
- Pseudowire—A bidirectional virtual circuit that, in the MPLS environment, is tunneled through the MPLS backbone and carried over a pair of LSPs.
- Destination Attachment Circuit—The circuit or virtual circuit at the other end of the Pseudowire that leads to the destination CE device.

Figure 4 shows the pathway of a VPWS between a source CE router and a PE router via an Attachment Circuit through a Pseudowire. In this case, the PE router serves as a virtual circuit switch. The “switching” mechanism it uses is a virtual entity called a Forwarder, which maps the Attachment Circuits to Pseudowires on a one-to-one basis. The MPLS Manager discovers both Forwarders and Pseudowires.

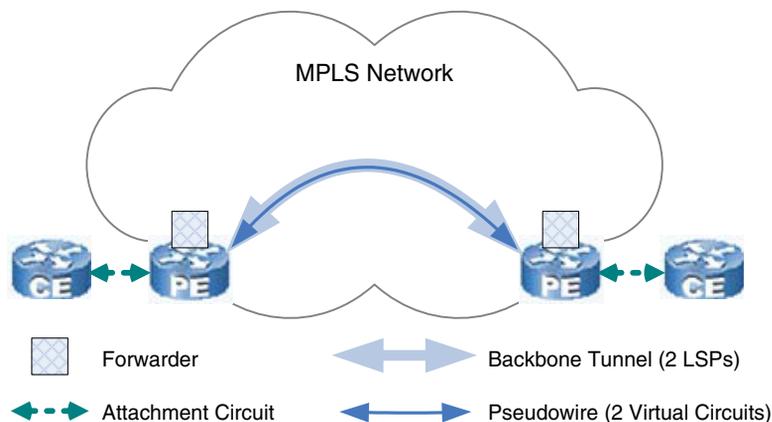


Figure 4: The L2VPN Domain Discovered by the MPLS Manager

The MPLS Manager discovers and manages the Forwarders and the Pseudowires for a VPWS and associates them to a VPN. In addition, the MPLS Manager associates the availability of a Pseudowire with the availability of a Pseudowire’s underlying end-to-end LSPs.

A Pseudowire is layered over an LSP, and for discovered Martini-implemented L2VPNs, is layered over the targeted LdpAdjacency session that terminates at the LdpProtocolEndpoints. LdpProtocolEndpoints reside on the source and destination PE devices. An LdpProtocolEndpoint represents the LDP entity on the PE device that exchanges the MPLS labels for the L2VPN. Figure 5 shows the LdpProtocolEndpoint and the path over which the LdpAdjacency session runs.

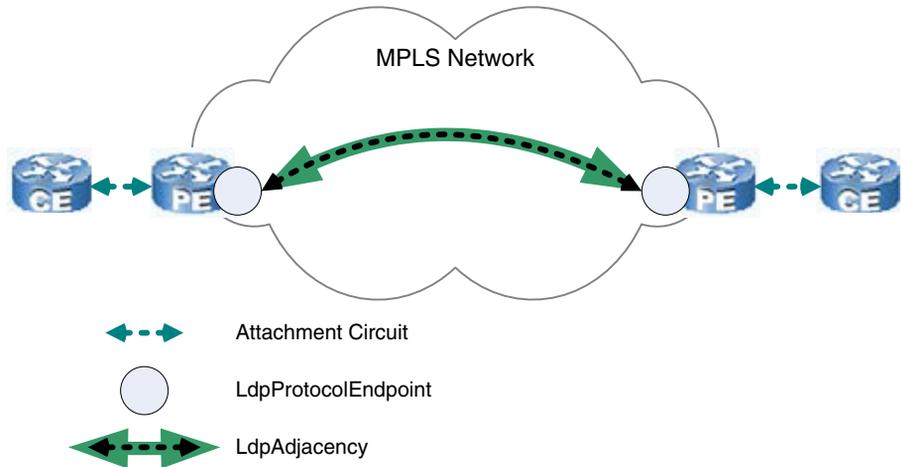


Figure 5: The L2VPN Domain—An LdpAdjacency Session

Transparent LAN Service

For Ethernet-based VLANs transported over a Pseudowire, the MPLS Manager provides the Transparent LAN Service (TLS) feature, which monitors end-to-end connectivity over the MPLS backbone. Figure 6 shows the relationship between the VLAN, the CEs, the Forwarders, and the Pseudowire that carries that VLAN across the MPLS backbone. The TLS feature enables the MPLS Manager to associate failures in the Pseudowire to failures of the VLAN.

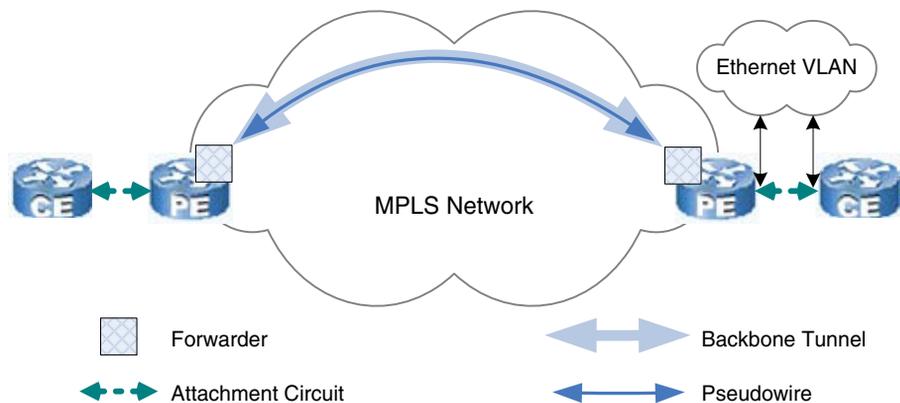


Figure 6: An Ethernet VLAN Associated with a Pseudowire

Remote Ping Functionality

Remote ping functionality, shown in Figure 7, allows the MPLS Manager to request periodic or on-demand remote pings from the following elements:

- PE router to PE router
- CE router to CE router
- PE to local CE
- PE to remote CE
- PE to unmanaged CE
- PE to VRF

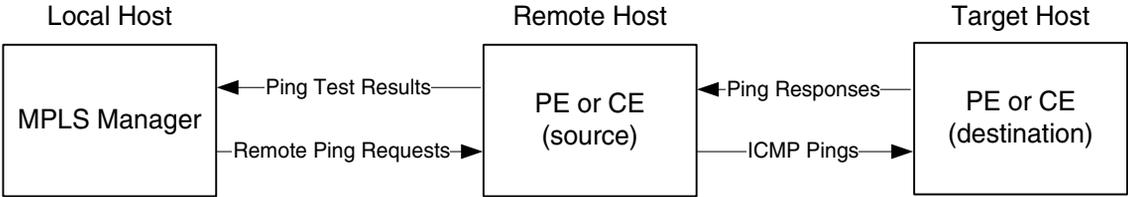


Figure 7: Remote Ping Functionality for the MPLS Manager

The MPLS Manager’s remote ping functionality uses Internet Control Message Protocol (ICMP) Echo requests to get information about the availability of network elements in the managed network. In most cases, a network element such as a router will respond to an ICMP Echo request if it is operationally up. Besides serving as an indicator of the operational status of an element, remote pings indicate whether a source element is able to reach the destination element.

Remote ping functionality can be automated or can be triggered as a server tool from the Domain Manager Administration Console view of the Global Console. For information about configuring remote ping parameters, see [Remote Ping Functionality](#) on page 97.

Monitoring

After the initial discovery of the MPLS and VPN elements, the MPLS Manager continuously monitors the status of the VRF (L3VPN), Forwarder (L2VPN), and LdpProtocolEndpoint (L2VPN) elements, and the remote ping results, by periodically sending SNMP polls to the PE routers in the managed MPLS environment. The results of the polling, in addition to the root-cause failure events received from Availability Manager, serve as input to the MPLS Manager correlation analysis. For information about SNMP polling, see [Customizing Groups and Settings](#) on page 69.

Analysis

The MPLS Manager detects configuration and other types of errors that occur when deploying and maintaining MPLS and VPN networks. It also correlates root-cause failures in the transport domain, diagnosed by Availability Manager, to impairments in the MPLS and VPN domains. As an example of this latter capability, also known as cross-domain impact correlation analysis, the MPLS Manager indicates the LSPs and corresponding VRFs and VPNs that have been impaired, or impacted, by an underlying PE router failure.

Notifications

The MPLS Manager sends the results of its root-cause and impact analysis to the Global Manager in the form of notifications, which are displayed in the Notification Log Console view of the Global Console. The root-cause and impact events notified by the MPLS Manager identify the MPLS Manager by its domain name in the Source attribute. Users can double-click a notification to view detailed information about the notification.

For information about the notifications created by the MPLS Manager and reported to the Global Manager, see [MPLS and VPN Elements and Their Failures](#) on page 9 and [MPLS Cross-Domain Impact Correlation Analysis](#) on page 35. For information about viewing the notifications, see [Viewing MPLS Notifications, Maps, and Containment](#) on page 49. For a list of notifications related to the underlying network, see [Root-Cause Notifications from Availability Manager](#) on page 125.

2

MPLS and VPN Elements and Their Failures

This chapter describes the MPLS, L3VPN, and L2VPN elements discovered and managed by the MPLS Manager and the various events notified for the elements.

In addition, it includes descriptions of MPLS, L3VPN, and L2VPN element attributes.

The MPLS Manager identifies configuration and other errors that occur when deploying and maintaining MPLS and VPN networks, and reports the errors to the Global Manager. It also:

- Correlates MPLS, L3VPN, and L2VPN impairments with transport root-cause failures received from the Availability Manager to identify MPLS, L3VPN, and L2VPN impacts, as explained in [MPLS Cross-Domain Impact Correlation Analysis](#) on page 35, and reports the impacts to the Global Manager.
- Performs remote ping analysis to identify L3VPN impacts, as explained in [Remote Ping Functionality](#) on page 97, and reports the impacts to the Global Manager.

MPLS and VPN Elements and Relationships

The MPLS Manager builds a data model of the discovered MPLS, L3VPN, and L2VPN elements in the managed MPLS environment. The model represents the discovered elements as instances of the InCharge Common Information Model (ICIM) classes.

The MPLS core elements are represented by instances of the following ICIM classes:

- MPLSService
- LSP
- LSPHop
- LSPInSegment
- LSPOutSegment

The L3VPN elements are represented by instances of the following ICIM classes:

- VPN
- VRF
- RouteTarget

The L2VPN elements are represented by instances of the following ICIM classes:

- VPN
- Forwarder
- PseudoWire
- LdpProtocolEndpoint
- LdpAdjacency

During the discovery post-processing phase, MPLS Manager creates the relationships and connections between the MPLS and VPN elements. Typically, every relationship has an inverse relationship. For example, the relationship PartOf is the inverse relationship of ComposedOf.

For an illustration of the relationships and connections between the MPLS and VPN elements, see the *EMC Smarts MPLS Manager Discovery Guide Supplement*. For general descriptions of ICIM classes, relationships, and connections, see the *EMC Smarts ICIM Reference*.

Summary of Events for L3VPN Networks

The MPLS Manager creates an event notification for each error that it detects in a managed L3VPN network. Notifications are displayed in the Global Console.

Table 3 lists the events notified by the MPLS Manager for a managed L3VPN network, including the condition for each event. The table also identifies the managed elements for which the events are notified.

Table 3: L3VPN Events Notified by the MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
VRF	Down	The VRF is operationally down: Either the VRF has no associated interfaces or it has one or more associated interfaces and all of them are down.*
	NoRoutes	The VRF has no routes in its routing table: The VRF has one or more associated interfaces but all of them are unnumbered. (An unnumbered interface has no IP address assigned to it.)
	WarningThresholdCrossed	The mid-threshold number of VRF routes has been crossed for this VRF.
	MaxRoutesReached	The maximum number of VRF routes has been reached for this VRF.
RouteTarget	Misconfiguration	A route target has been configured but is not being used by any of the VRFs in the managed MPLS environment.
* MPLS Manager also notifies VRF impacts for each of the down interfaces.		

Summary of Events for L2VPN Networks

The MPLS Manager also creates an event notification for each error that it detects in a managed L2VPN network. Notifications are displayed in the Global Console.

Table 4 lists the events notified by the MPLS Manager for a managed L2VPN network, including the condition for each event. The table also identifies the managed elements for which the events are notified.

Table 4: L2VPN Events Notified by the MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
Forwarder	IsDown	The operational or administrative status of the Forwarder is down. Note that this event is implemented only for Juniper devices.
PseudoWire	Disconnected	The Pseudowire is disconnected because the operational or administrative status of the endpoint Forwarders is down. Note that this event is implemented only for Juniper devices.
	Down	The Pseudowire is down because at least one of the corresponding Forwarders is down, or there is an underlying physical transport problem. Note that this event is implemented only for Juniper devices.
LdpProtocolEndpoint	IsDown	The LdpProtocolEndpoint has been impaired by failures of the corresponding terminating device or of the corresponding peer device. Note that this event is implemented only for Cisco devices.
LdpAdjacency	Disconnected	The LdpAdjacency has been impaired by failure of at least one of the corresponding LdpProtocolEndpoints. Note that this event is implemented only for Cisco devices.
	Down	The LdpAdjacency is down because at least one of the corresponding LdpProtocolEndpoints is down, or there is an underlying physical transport problem. Note that this event is implemented only for Cisco devices.

MPLS Core Elements

This section describes the discovered elements specific to MPLS networks:

- MPLSService
- LSP
- LSPHop
- LSPInSegment
- LSPOutSegment

MPLSService

An MPLS service is a logical element created for each router (PE, P, CE) discovered in the managed MPLS environment, even if the router (CE router, for example) does not support MPLS. For definitions of P, PE, and CE routers, see [MPLS Terminology](#) on page 113.

The relationships created for an MPLS service depend on the type of device (PE, P, CE) hosting the MPLS service and the type of VPN (L3VPN, L2VPN) supported by the MPLS service. For a description of the relationships created for an MPLS service, see the *EMC Smarts MPLS Manager Discovery Guide Supplement*.

Attributes for MPLSService

Table 5 lists key attributes for MPLSService.

Table 5: Attributes for MPLSService

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
DeviceType	Type of router hosting this MPLS service.	<ul style="list-style-type: none"> • PE • P • CE • NON_MPLS • Other
Name	Name assigned to this MPLS service. The name format is MPLS- <i><dev></i> , where <i><dev></i> is the name or IP address of the router hosting this MPLS service.	String

Table 5: Attributes for MPLSService (continued)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
NumberOfVRFs	Number of VRFs maintained by the router hosting this MPLS service. Applicable to PE routers only; otherwise, set to 0.	Integer
Supports_LSR_MIB	True if the router hosting this MPLS service supports the SNMP MPLS-LSR MIB.	Boolean: true or false
Supports_VPN_MIB	True if the router hosting this MPLS service supports the SNMP MPLS-VPN MIB.	Boolean: true or false
TotalVRFRoutes	Total number of VPN routes in the VRFs maintained by the router hosting this MPLS service. Applicable to PE routers only; otherwise, set to 0.	Integer

LSP

An LSP is a fixed data-forwarding path traversed by labeled packets through an MPLS network. An LSP starts at one PE router and ends at another PE router, and consists of a sequence of LSP hops in which a packet travels from router to router via a label switching mechanism.

An LSP can be established dynamically, based on normal routing mechanisms, or through configuration. Once an LSP is established, all subsequent packets follow the same path.

The MPLS Manager discovers LSPs in the context of VPNs. Thus, for L3VPNs, the MPLS Manager discovers only those LSPs between PEs that have VPN routes configured between them; for L2VPNs, the MPLS Manager discovers only those LSPs between PEs that have Pseudowires configured between them. In situations where not all of the routers in the MPLS network are managed by the MPLS Manager, a discovered LSP may represent something less than the entire LSP path.

Attributes for LSP

Table 6 lists key attributes for LSP.

Table 6: Attributes for LSP

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
LSPId	A 32-bit integer that uniquely identifies this LSP within the scope of the managed MPLS environment. Its value is the LSP's destination subnet, which is an IP address on the destination PE router for this LSP. For an LSP underlying an L3VPN, the IP address is typically that of a BGP speaker.	String: IpAddress, an application-wide type representing a 32-bit internet address
Name	Name assigned to this LSP. The name format is <code>LSP-<dev1>-><dev2></code> , where: <ul style="list-style-type: none"> <code><dev1></code> is the name or IP address of the source PE router for this LSP. <code><dev2></code> is the name or IP address of the destination PE router for this LSP. 	String

LSPHop

An LSP hop is a unidirectional logical link between two routers in an MPLS network across which MPLS-labeled packets are sent. No label processing occurs over the logical link.

An exception to this definition is the last hop of an LSP, across which the packets may be unlabeled due to *penultimate hop popping* (also known as penultimate label popping)—see definition in [MPLS Terminology](#) on page 113. In this case, the Label attribute of the LSP hop is 3, although the packets are, in fact, unlabeled.

Note: For VPN packets and penultimate hop popping, the packets retain their inner label when traversing the last hop of an LSP.

Attributes for LSPHop

Table 7 lists key attributes for LSPHop.

Table 7: Attributes for LSPHop

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Label	Label assigned to the MPLS packets traversing this LSP hop. The label is equal to the label assigned to the LSP outsegment to which this LSP hop connects. For the last hop of an LSP for which penultimate hop popping is in effect, the Label is set to 3—the implicit Null label.	Integer: in the range 0 through 1048575
LSPId	LSP identifier for the LSP to which this LSP hop belongs.	String: IpAddress, an application-wide type representing a 32-bit internet address
Name	Name assigned to this LSP hop. The name format is LSPHop- <label>/<destdev>/<destif>, where: <ul style="list-style-type: none"> • <label> is the label for this LSP hop. • <destdev> is the name or IP address of the destination router for this LSP hop. • <destif> is the interface number of the incoming interface associated with this LSP hop on the destination router. For an LSP hop having Label = 3, the name format is LSPHop-POP/<srcdev>/<srcif>-<key>, where: <ul style="list-style-type: none"> • <srcdev> is the name or IP address of the source router for this LSP hop. • <srcif> is the interface number of the outgoing interface associated with the LSP outsegment to which this LSP hop connects. • <key> is a value obtained from the MPLS-LSR MIB that uniquely identifies the LSP outsegment to which this LSP hop connects. Note that the <key> value is used to distinguish between different LSP hops that are all POP (Label = 3), originate on the same router, and use the same outgoing interface. 	String

LSPInSegment

An LSP insegment is an incoming label in the MPLS forwarding table of a PE or P router. An MPLS forwarding table maps LSP insegments (incoming labels) to LSP outsegments (outgoing labels and associated outgoing interfaces.) Each LSP insegment and LSP outsegment pair represents an entry in the MPLS forwarding table.

For an ingress PE router, LSP insegments have no meaning and therefore are not created by the MPLS Manager. An LSP starts at, and is determined by, the LSP outsegment chosen by the ingress PE router.

Attributes for LSPInSegment

Table 8 lists key attributes for LSPInSegment.

Table 8: Attributes for LSPInSegment

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Label	Label assigned to this LSP insegment.	Integer: in the range 0 through 1048575
LSPId	LSP identifier for the LSP hop to which this LSP insegment connects.	String: IpAddress, an application-wide type representing a 32-bit internet address
Name	Name assigned to this LSP insegment. The name format is LSPInSegment - <i><label>/<dev></i> , where: <ul style="list-style-type: none"> <i><label></i> is the label assigned to this LSP insegment. <i><dev></i> is the name or IP address of the router hosting the MPLS forwarding table to which this LSP insegment belongs. 	String

LSPOutSegment

An LSP outsegment is an outgoing label in the MPLS forwarding table of a PE or P router. By means of the MPLS forwarding table, a labeled packet coming into a router is relabeled with the appropriate outgoing label and sent out over the appropriate outgoing interface.

For an egress PE router, LSP outsegments have no meaning and therefore are not created by MPLS Manager: An LSP ends at the LSP insegment of the egress PE router, or if penultimate hop popping is in effect, ends at the last LSP hop for the LSP.

Attributes for LSPOutSegment

Table 9 lists key attributes for LSPOutSegment.

Table 9: Attributes for LSPOutSegment

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
ifIndex	Interface number of the outgoing interface associated with this LSP outsegment.	Integer
Key	A value obtained from the MPLS-LSR MIB that uniquely identifies this LSP outsegment. Note that the Key attribute is used to distinguish between different LSP outsegments that are all POP (Label = 3), reside on the same router, and use the same outgoing interface. For clarification, see the Label and Name attribute descriptions in this table.	String
Label	Label assigned to this LSP outsegment. If this LSP outsegment connects to the last hop of an LSP for which penultimate hop popping is in effect, the Label is set to 3 (implicit Null label).	Integer: in the range 0 through 1048575
LSPId	LSP identifier for the LSP hop to which this LSP outsegment connects.	String: IpAddress, an application-wide type representing a 32-bit internet address

Table 9: Attributes for LSPOutSegment (continued)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Name	<p>Name assigned to this LSP outsegment. The name format is LSPOutSegment- <label>/<dev>/<ifindex>, where:</p> <ul style="list-style-type: none"> • <label> is the label assigned to this LSP outsegment. • <dev> is the name or IP address of the router hosting the MPLS forwarding table to which this LSP outsegment belongs. • <ifindex> is the ifIndex attribute value for this LSP outsegment. <p>For an LSP outsegment having Label = 3, the name format is LSPOutSegment-POP/<dev>/<ifindex>-<key>, where:</p> <ul style="list-style-type: none"> • <dev> is the name or IP address of the router hosting the MPLS forwarding table to which this LSP outsegment belongs. • <ifindex> is the ifIndex attribute value for this LSP outsegment. • <key> is the Key attribute value for this LSP outsegment. 	String
NextHopIP	IP address of the next router to receive the packets sent from this LSP outsegment.	String: IpAddress, an application-wide type representing a 32-bit internet address

L3 VPN Elements

This section describes the discovered elements specific to Layer 3 VPN networks configured and provisioned over MPLS networks:

- VPN (Layer 3)
- VRF
- RouteTarget

VPN (Layer 3)

A VPN in an L3VPN is a collection of VPN Routing and Forwarding (VRF) instances, configured on PE routers in the MPLS network, that are members of the same virtual private network. All of the functions associated with establishing, maintaining, and operating an L3VPN take place in the PE routers.

The P routers are not aware of the L3VPNs; they forward packets over the established LSPs. Similarly, the CE routers are not aware of the L3VPNs; they route IP packets in accordance with the customer's established addressing and routing schemes.

There are three types of VPN:

- Full mesh—Each customer site can communicate directly with every other customer site in the VPN.
- Hub and spoke—All traffic flows to/from a central hub site.
- Partial mesh—Some customer sites can communicate directly with other customer sites in the VPN. Essentially, a partial-mesh VPN is a hub-and-spoke VPN that has multiple hubs.

Attributes for VPN (Layer 3)

Table 10 lists key attributes for VPN in an L3VPN network.

Table 10: Attributes for VPN (Layer 3)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Name	<p>Name assigned to this VPN. For a full-mesh VPN in an L3VPN network, the name format is <code>VPN-<i><rt1></i></code>, where <i><rt1></i> is the name of the route target associated with this VPN.</p> <p>For a hub-and-spoke VPN in an L3VPN, the name format is <code>VPN-<i><rt2></i>:<i><rt3></i></code>, where:</p> <ul style="list-style-type: none"> • <i><rt2></i> is the name of the one route target associated with this VPN. • <i><rt3></i> is the name of the other route target associated with this VPN. 	String

Table 10: Attributes for VPN (Layer 3) (*continued*)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Topology	Topology of this VPN, which in an L3VPN network is determined by examining the route targets exported and imported by the VRFs comprising this VPN.	<ul style="list-style-type: none"> • FullMesh • Hub&Spoke
VPNType	Type of VPN. For a VPN in an L3VPN network, this attribute is set to RFC_2547.	<ul style="list-style-type: none"> • RFC_2547 • L2VPN

VRF

Although originally designed as a highly critical component in L3VPNs, the VRF has become an important component in Kompella-implemented L2VPNs. Both VPN architectures use Multiprotocol Border Gateway Protocol (MBGP) as their signaling protocol.

A VRF is a VPN Routing and Forwarding instance, maintained by a PE router, that contains the routing information defining a customer VPN site. A PE router maintains a VRF for each of its directly connected customer VPN sites. Multiple VRFs on multiple PE routers comprise a VPN.

A VRF consists of the following components:

- An IP routing table
- A derived VPN-specific forwarding table
- A set of PE router interfaces (tied to the locally attached customer VPN site) that use the forwarding table
- A set of rules and routing protocols that determine what goes into the forwarding table

The VRF stores packet forwarding information for the routes that are particular to the VPN to which the VRF belongs. Each route in the VRF is associated with two labels: an outer label used to route the packet through the MPLS network to the appropriate egress PE router, and an inner label used to deliver the packet to the correct VRF and correct end user.

It is interesting to note that because a PE router may have the same IP address on multiple interfaces, the Availability Manager source for MPLS Manager tags each of the IP addresses with a *route distinguisher* value that is unique to a particular VRF, to form unique VRF IP addresses. The route distinguisher is the means by which the PE router and the MPLS Manager keep track of overlapping customer IP address spaces.

Attributes for VRF

Table 11 lists key attributes for VRF. MPLS Manager uses these and other attributes to diagnose VRF configuration errors.

Table 11: Attributes for VRF

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
MaxRoutes	Denotes the maximum number of routes that this VRF is configured to hold.	Integer 0 signifies that the maximum route threshold is not set for this VRF.
MidRouteThreshold	Denotes the mid-level water marker for the number of routes that this VRF is configured to hold.	Integer 0 signifies that the mid-route threshold is not set for this VRF.
Name	Name assigned to this VRF. The name format is <code>VRF-<vrfname>/<dev></code> , where: <ul style="list-style-type: none"> • <code><vrfname></code> is the VRFName attribute value for this VRF. • <code><dev></code> is the name or IP address of the router hosting this VRF. 	String
NumberOfRoutes	Number of routes currently held by this VRF.	Integer
RouteDistinguisher	A value included in the network route advertisement for this VRF, to identify the VPN to which the route belongs.	String
VRFName	A value that distinguishes this VRF from other VRFs within the scope of the managed MPLS environment.	String

Events for VRF

Table 12 lists the events notified for VRF.

Table 12: Events for VRF

EVENT	DESCRIPTION
Down	This VRF is operationally down: Either the VRF has no associated interfaces or it has one or more associated interfaces and all of them are down.*
NoRoutes	This VRF has no routes in its routing table: Either the VRF has no associated interfaces or it has one or more associated interfaces and all of them are unnumbered. (An unnumbered interface has no IP address assigned to it.) The routes considered when computing this event include the advertised routes received from both the VRF's locally attached customer VPN site and the VRF's peer VRFs.
WarningThresholdCrossed	The number of routes held by this VRF exceeds the mid-route threshold (MidRouteThreshold attribute value) configured for this VRF. For this event to occur, the PE router hosting this VRF must support the MPLS-VPN MIB and must be SNMP-instrumented. MPLS Manager monitors the number of routes in the VRF and generates an event when the MidRouteThreshold attribute value is crossed.
MaxRoutesReached	The maximum number of routes held by this VRF equals the maximum route threshold (MaxRoutes attribute value) configured for this VRF. For this event to occur, the PE router hosting this VRF must support the MPLS-VPN MIB and must be SNMP-instrumented. MPLS Manager monitors the number of routes in the VRF and generates an event when the MaxRoutes attribute value is reached.
* MPLS Manager also notifies VRF impacts for each of the down interfaces.	

Note that the condition *VRF has no associated interfaces* triggers both a VRF Down event and a NoRoutes event but that the NoRoutes event is suppressed, meaning that the NoRoutes event is not reported to the Global Manager.

Also note that the VRF does not reject new routes even if the number of routes exceeds the maximum route threshold.

RouteTarget

Although originally designed as a key component in L3VPNs, the route target has also become a key component in Kompella-implemented L2VPNs. Both VPN architectures use MBGP as their signaling protocol.

A route target identifies a set of customer VPN sites to which a PE router distributes routes. It is used to set up peering relationships between the VRF instances that belong to the same VPN.

A VRF is configured with a route target export list and a route target import list. The host PE router inserts the VRF's export list into route advertisements for the VRF, and accepts route advertisements having at least one route target matching a member of the VRF's import list.

Attributes for RouteTarget

Table 13 lists key attributes for RouteTarget.

Table 13: Attributes for RouteTarget

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Key	Value of this route target; for example, 100:3000.	String: a 64-bit quantity
Name	Name assigned to this route target. The name format is RT- <i><key></i> , where: <i><key></i> is the Key attribute value for this route target.	String

Events for RouteTarget

A single event is notified for RouteTarget and is called Misconfiguration. The event indicates that a route target has been configured but is not being used by any of the VRFs in the managed MPLS environment.

L2VPN Elements

This section describes the discovered elements specific to Layer 2 VPN networks configured and provisioned over MPLS networks:

- VPN (Layer 2)
- Forwarder
- PseudoWire
- LdpProtocolEndpoint
- LdpAdjacency

VPN (Layer 2)

A VPN in an L2VPN network is a collection of Forwarder and Pseudowire instances (and, for Kompella-implemented L2VPNs, VRFs and route targets), configured on PE routers in the MPLS network, that are members of the same virtual private network. All of the functions associated with establishing, maintaining, and operating an L2VPN take place in the PE routers.

The P routers are not aware of the L2VPNs; they forward packets over the established LSPs. Similarly, the CE devices, which may be routers, switches, hosts, or just about anything that the L2VPN customer wants to connect to an L2VPN, are not aware of the L2VPNs; they operate without any knowledge of the existence of L2VPNs.

Note: In an L2VPN, the CE is attached to the PE via an Attachment Circuit, which may be a physical or logical link.

The MPLS Manager discovers and creates the following type of VPN for Martini-implemented L2VPNs: point-to-point full mesh. Because the signaling protocol (LDP) used by Martini-implemented L2VPNs does *not* advertise L2VPN membership information among the PE routers, the MPLS Manager cannot determine to which L2VPN a discovered Pseudowire belongs. Accordingly, the MPLS Manager considers each discovered Pseudowire (Martini Tunnel) a point-to-point full-mesh L2VPN.

The MPLS Manager creates a VPN element (object) for each discovered Martini-implemented L2VPN. If you wish to change this behavior so that the MPLS Manager does not create VPN elements for the discovered L2VPNs, set the `L2VPN_CREATE_VPN` parameter to `FALSE` in the MPLS Manager's `LOCAL.import` file, as explained in the *EMC Smarts MPLS Manager Configuration Guide*.

The MPLS Manager discovers and creates the following types of VPN for Kompella-implemented L2VPNs: full mesh and hub and spoke. Because the signaling protocol (MBGP) used by Kompella-implemented L2VPNs *does* advertise L2VPN membership information among the PE routers, the MPLS Manager can determine to which L2VPN a discovered Pseudowire belongs.

Attributes for VPN (Layer 2)

Table 14 lists key attributes for VPN in an L2VPN network.

Table 14: Attributes for VPN (Layer 2)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Name	Name assigned to this VPN. For a Martini-implemented L2VPN, the name format is <code>VPN-L2 - <pseudowire></code> , where <code><pseudowire></code> is the name of the Pseudowire instance associated with this VPN. For a Kompella-implemented L2VPN, the name format is the same as for a Layer 3 VPN; for details, see Attributes for VPN (Layer 3) on page 20.	String
Topology	Topology of this VPN. For a Martini-implemented L2VPN, this attribute is set to FullMesh. For a Kompella-implemented L2VPN, this attribute is determined by examining the route targets exported and imported by the VRFs comprising this VPN.	<ul style="list-style-type: none"> • FullMesh • Hub&Spoke
VPNType	Type of VPN. For a VPN in an L2VPN network, this attribute is set to L2VPN.	<ul style="list-style-type: none"> • RFC_2547 • L2VPN

Forwarder

A Forwarder is the logical entity within a PE router that makes switching and forwarding decisions. It connects a customer-side Attachment Circuit—a VLAN or an Ethernet port, for example—to an MPLS-side Pseudowire—a Martini Tunnel, for example. If the Attachment Circuit is a VLAN, the Forwarder is LayeredOver the VLAN in the MPLS Manager data model.

- The Attachment Circuit that connects the source PE to the source CE
- The Pseudowire that connects the source PE to the destination PE
- The Attachment Circuit that connects the destination PE to the destination CE

Attributes for Forwarder

Table 15 lists key attributes for Forwarder.

Table 15: Attributes for Forwarder

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
OutSegment	Virtual Connection Identifier of the Pseudowire that is associated with this Forwarder.	Integer
VPNType	Type of VPN to which this Forwarder belongs.	<ul style="list-style-type: none"> • UNKNOWN • OTHER,BGP_IP_VPN • BGP_L2_VPN • L2_CIRCUIT • OPTICAL_VPN • VP_OXC, CCC
VPNName	Name of the VPN to which this Forwarder belongs.	String
Index	Index of this Forwarder in the MIB.	Integer
TunnelName	Name of the tunnel that terminates on this Forwarder.	String
TunnelType	Type of tunnel that terminates on this Forwarder.	String
LocalSiteID	Identifier of the site to which this Forwarder belongs.	Integer
RemoteSiteID	Identifier of the site to which this Forwarder's peer belongs.	Integer
Status_CLI	Denotes whether a Forwarder is operational or not.	String
VC_ID	Virtual Connection Identifier of the Pseudowire associated with this Forwarder.	Integer

Table 15: Attributes for Forwarder *(continued)*

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Trans_Demux	Outgoing Virtual Connection Identifier of the Pseudowire associated with this Forwarder.	Integer
Rcv_Demux	Incoming Virtual Connection Identifier of the Pseudowire associated with this Forwarder.	Integer
PeerAddress	IP Address for the peer Forwarder.	String: IpAddress, an application-wide type representing a 32-bit internet address
VLANs	List of VLANs underlying this Forwarder. (A Forwarder is LayeredOver a VLAN.)	String

Events for Forwarder

A single event is notified for Forwarder and is called IsDown. The event indicates that the operational or administrative status of the Forwarder is down. Note that this event is implemented only for Juniper devices.

PseudoWire

A Pseudowire is a bidirectional virtual connection between two PE routers. The MPLS Manager creates the Pseudowires during the discovery post-processing phase, when it also creates the relationships and connections between the L2VPN and discovered MPLS core elements (LSPs and LSPHops).

The Pseudowire traverses the MPLS network through a secure "tunnel."

Attributes for PseudoWire

Table 16 lists key attributes for PseudoWire.

Table 16: Attributes for PseudoWire

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Address1	IP Address for Termination Point 1.	String: IpAddress, an application-wide type representing a 32-bit internet address
Address2	IP Address for Termination Point 2.	String: IpAddress, an application-wide type representing a 32-bit internet address
VC_ID	Virtual Connection Identifier of the Pseudowire.	Integer
Demux1	Incoming Virtual Connection Identifier of the Pseudowire.	Integer
Demux2	Outgoing Virtual Connection Identifier of the Pseudowire.	Integer
Termination1DisplayName	Display name of the neighboring Termination Point 1.	String
Termination2DisplayName	Display name of the neighboring Termination Point 2.	String
IsFullyConnected	Indicates whether this tunnel is fully instrumented on both of its termination points.	Boolean: true or false
Status1	Status of this tunnel as reported by Termination Point 1.	<ul style="list-style-type: none"> • UNKNOWN • DOWN • UP
Status2	Status of this tunnel as reported by Termination Point 2.	<ul style="list-style-type: none"> • UNKNOWN • DOWN • UP

Events for PseudoWire

Table 17 lists the events notified for PseudoWire.

Table 17: Events for PseudoWire

EVENT	DESCRIPTION
Disconnected	The Pseudowire is disconnected because the operational or administrative status of the endpoint Forwarders is down. Note that this event is implemented only for Juniper devices.
Down	The Pseudowire is down because at least one of the corresponding Forwarders is down, or there is an underlying physical transport problem. Note that this event is implemented only for Juniper devices.

LdpProtocolEndpoint

The LdpProtocolEndpoint represents the Label Distribution Protocol (LDP) entity on a PE router that is responsible for exchanging the MPLS labels for a Martini-implemented L2VPN. Each of these protocol endpoints on a PE router represents a targeted LDP session to the L2VPN peer PE.

Attributes for LdpProtocolEndpoint

Table 18 lists key attributes for LdpProtocolEndpoint.

Table 18: Attributes for LdpProtocolEndpoint

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Index	Index of this LDP Protocol Endpoint in the MIB.	Integer
IsTargetedPeer	Indicates that this protocol termination point is to a targeted peer.	boolean
PeerAddressType	Type of the internetwork layer address used for Extended Discovery. This object indicates how the value of mplsLdpEntityTargetedPeerAddr is to be interpreted.	Integer
LocalAddress	IP Address for this protocol endpoint.	String: IpAddress, an application-wide type representing a 32-bit internet address

Table 18: Attributes for LdpProtocolEndpoint (continued)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
PeerAddress	IP Address for the peer of this protocol endpoint.	String: IpAddress, an application-wide type representing a 32-bit internet address
AdminStatus	Administration status of the protocol endpoint.	<ul style="list-style-type: none"> • ENABLE • DISABLE • UNKNOWN
OperStatus	Operational status of the protocol endpoint.	<ul style="list-style-type: none"> • ENABLE • DISABLE • UNKNOWN
AdjacencyHoldTime	LifeTime of the protocol endpoint in seconds.	Integer
AdjacencyHoldTimeRemaining	LifeTime of the protocol endpoint in seconds.	Integer
HostDescription	Description of hosting system.	String
AdjacencyEstablished	Indicates that the LDP Adjacency is established.	Boolean: true or false

Events for LdpProtocolEndpoint

A single event is notified for LdpProtocolEndpoint and is called IsDown. The event indicates that the LdpProtocolEndpoint has been impaired by failures of the corresponding terminating device or of the corresponding peer device. Note that this event is implemented only for Cisco devices.

LdpAdjacency

The LDPAdjacency represents a targeted LDP session between two peer PE routers involved in a Martini-implemented L2VPN. This session is used for exchanging the Virtual Circuit ID for the Pseudowire between the two PE routers.

Attributes for LdpAdjacency

Table 19 lists key attributes for LdpAdjacency.

Table 19: Attributes for LdpAdjacency

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Endpoint1DisplayName	Display name of neighboring Endpoint 1.	String
Endpoint2DisplayName	Display name of neighboring Endpoint 2.	String
OperStatus1	Operational status of the protocol Endpoint 1.	<ul style="list-style-type: none"> • UNKNOWN • DOWN • UP
OperStatus2	Operational status of the protocol Endpoint 2.	<ul style="list-style-type: none"> • UNKNOWN • DOWN • UP
Established	Indicates whether ALL neighbors are connected to this link.	Boolean: true or false

Events for LdpAdjacency

Table 20 lists the events notified for LdpAdjacency.

Table 20: Events for LdpAdjacency

EVENT	DESCRIPTION
Disconnected	The LdpAdjacency has been impaired by failure of at least one of the corresponding LdpProtocolEndpoints. Note that this event is implemented only for Cisco devices.
Down	The LdpAdjacency is down because at least one of the corresponding LdpProtocolEndpoints is down, or there is an underlying physical transport problem. Note that this event is implemented only for Cisco devices.

Provisioning System Adapter Events

The MPLS Manager architecture may contain a specialized adapter that not only provides additional customer information about the provisioned VPNs, but also synchronizes VPN and customer information between MPLS Manager and the customer provisioning system. The adapter is located between the provisioning system and the Global Manager, and communicates with MPLS Manager and Availability Manager through the Global Manager.

Currently, only one such adapter is available: EMC Smarts Adapter for Cisco ISC (ISC Adapter), which interfaces with the Cisco Internet Solutions Center (ISC) provisioning system.

When the ISC Adapter cannot reconcile VPN provisioning data differences between the MPLS Manager and the provisioning system, it generates certain event notifications and reports them to the Global Manager. For a description of these event notifications, see the *EMC Smarts Adapter for Cisco ISC User's Guide*.

Underlying Transport Network Failures

To understand how the Availability Manager discovers and monitors the underlying transport network elements and diagnoses connectivity failures between those elements, see the *EMC Smarts IP Availability Manager User's Guide*.

To understand how the MPLS Manager correlates the underlying transport network failures received from the Availability Manager with MPLS symptoms to identify MPLS impacts, see [MPLS Cross-Domain Impact Correlation Analysis](#) on page 35.

MPLS Cross-Domain Impact Correlation Analysis

The MPLS Manager calculates the impacts on MPLS, L3VPN, and L2VPN elements caused by underlying transport network failures, and reports the impacts to the Global Manager.

Impact Analysis Overview

The MPLS Manager receives transport root-cause problem events from the Availability Manager, including Router Down, Interface Down, Interface Disabled, NetworkConnection Down. For a complete list of problem events received by the MPLS Manager from the Availability Manager, see [Root-Cause Notifications from Availability Manager](#) on page 125.

- When the MPLS Manager receives a transport problem event for a managed L3VPN network, it correlates the problem with VPN, VRF, and LSP impairments and reports the impairments as impact notifications to the Global Manager.
- When the MPLS Manager receives a transport problem event for a managed L2VPN network, it correlates the problem with Forwarder, PseudoWire, LdpProtocolEndpoint, and LSP impairments and reports the impairments as impact notifications to the Global Manager.

The Global Manager responds by adding the impacts to the transport root-cause problem notification received from the Availability Manager.

Figure 8 shows the flow of information between the components in an MPLS Manager deployment to achieve MPLS and global impact analysis.

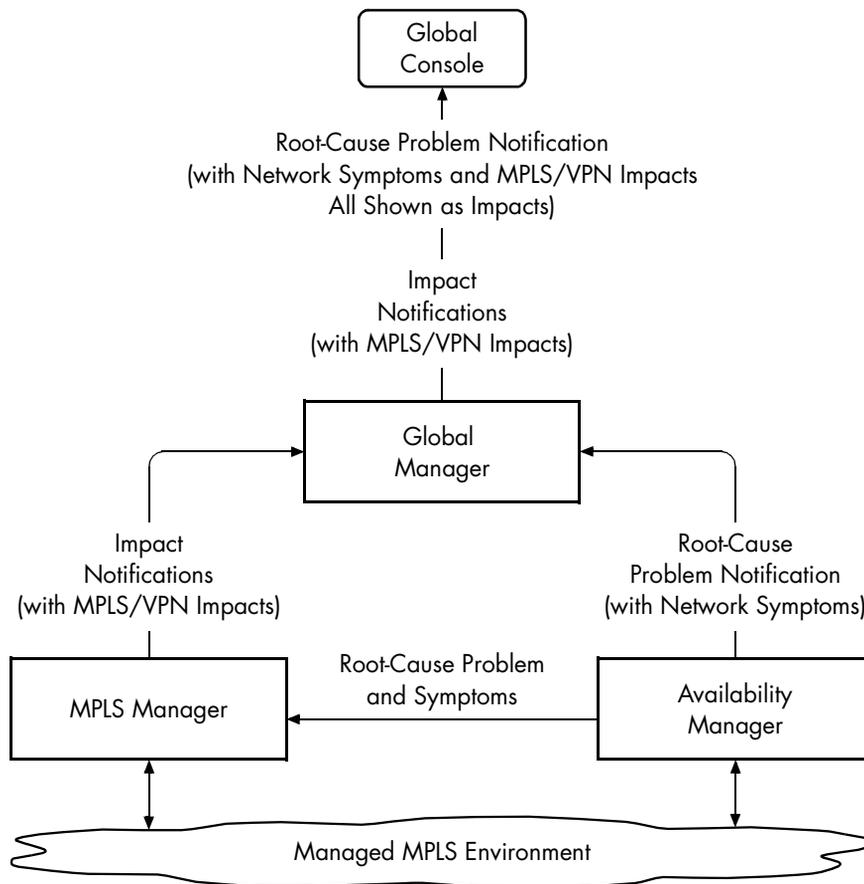


Figure 8: MPLS and Global Impact Analysis

Impact Analysis Models

Figure 9 demonstrates how the MPLS Manager models (represents) the MPLS and L3VPN elements for a managed L3VPN network, and Figure 10 and Figure 11 demonstrate how the MPLS Manager models the MPLS and L2VPN elements for a managed L2VPN network. The underlying transport network elements in the models, shown as white text on black background, are managed by the Availability Manager. As such, the MPLS Manager receives the status of these elements from the Availability Manager.

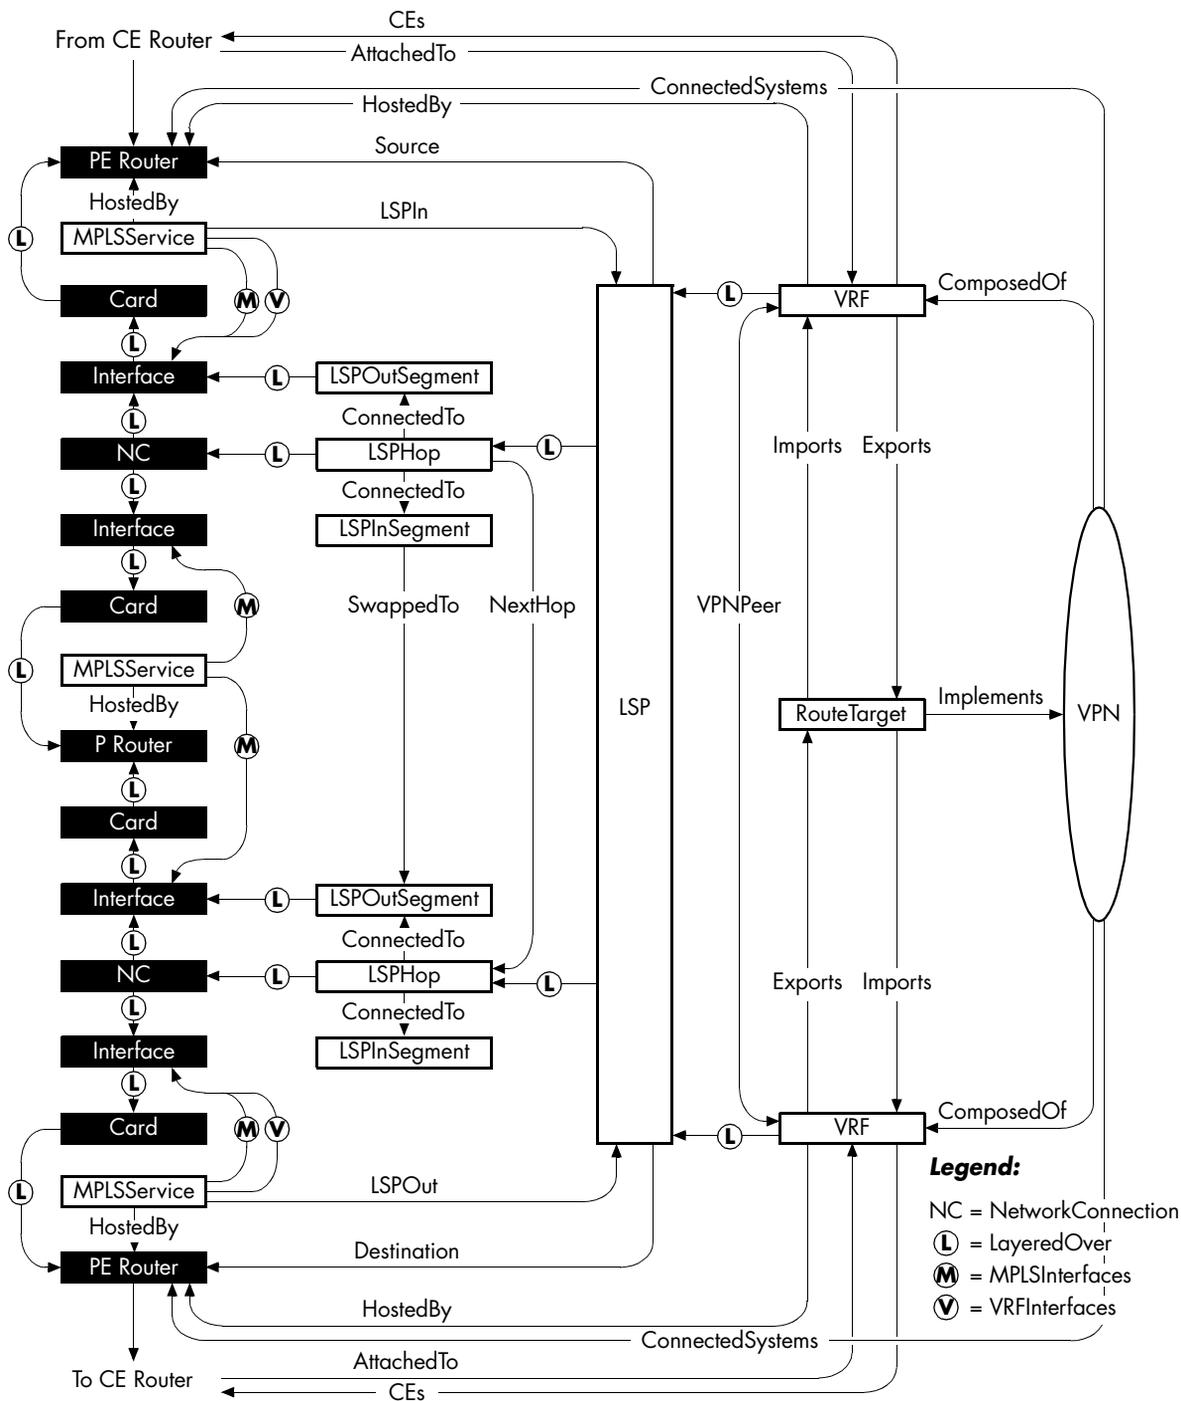


Figure 9: MPLS Manager View of a Managed L3VPN Network

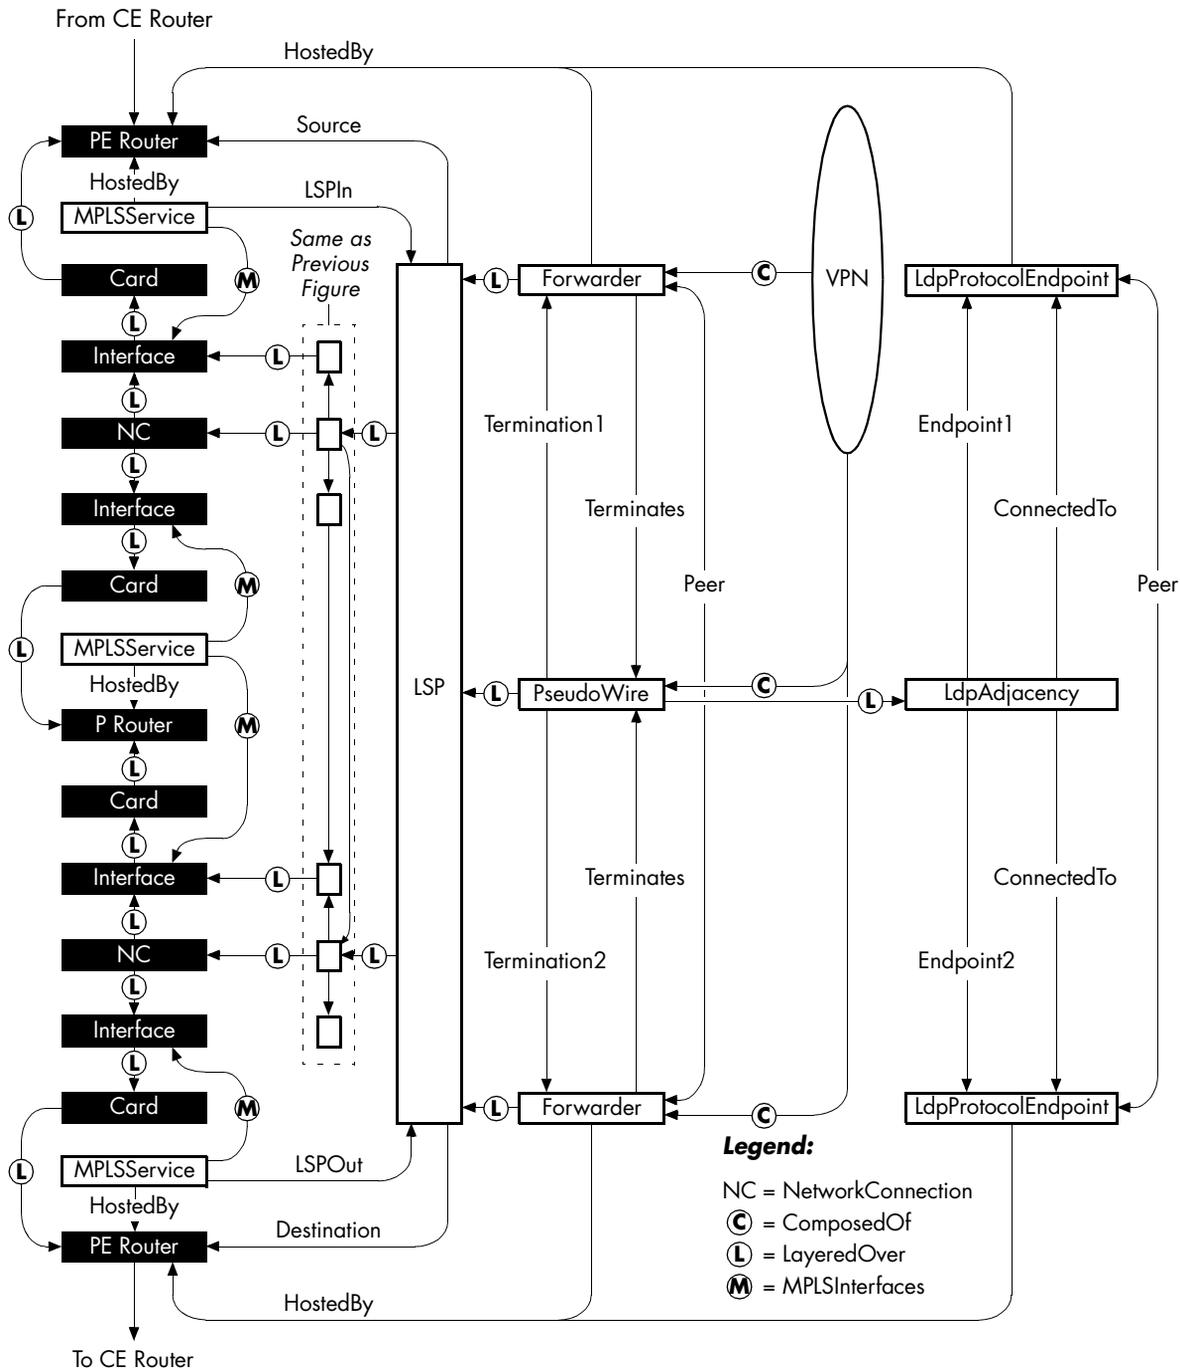


Figure 10: MPLS Manager View of a Managed Martini-Implemented L2VPN Network

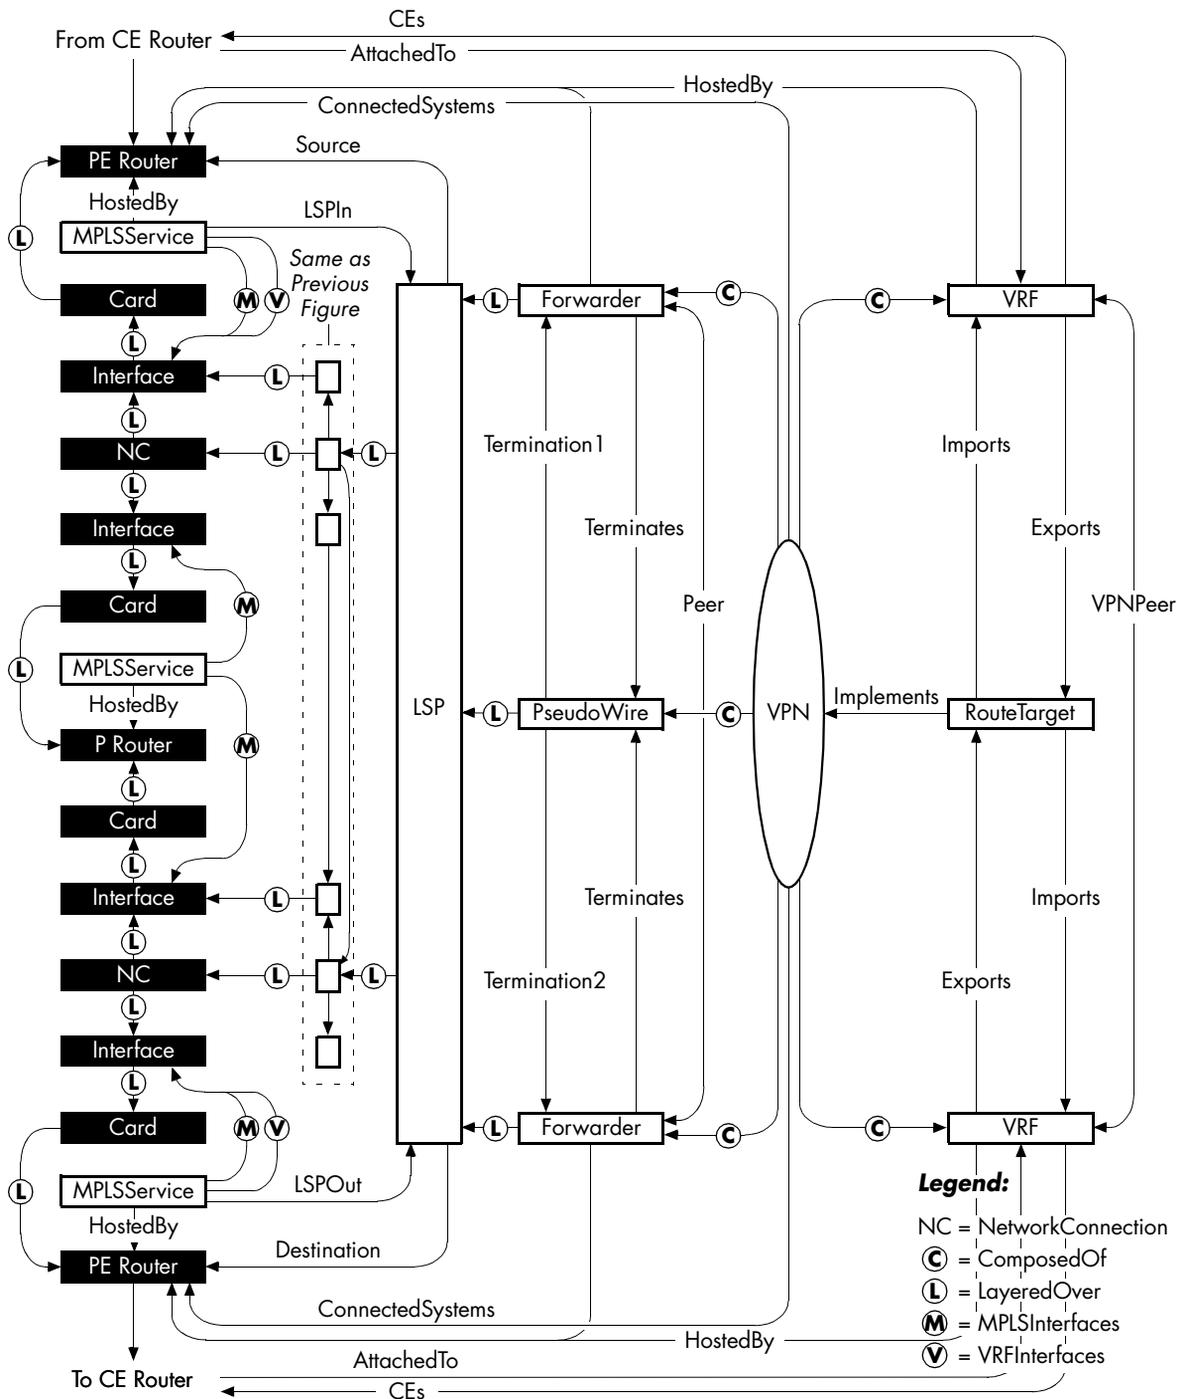


Figure 11: MPLS Manager View of a Managed Kompella-Implemented L2VPN Network

Impact Analysis Events

The MPLS Manager creates an impact event notification for each calculated impact. Notifications are imported by the Global Manager and displayed in the Global Console.

L3VPN Domain Impact Events

Table 21 lists the impact events notified by the MPLS Manager for the L3VPN domain, including the condition for each event. The table also identifies the managed elements for which the impact events are notified.

Table 21: L3VPN Impact Events Notified by the MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
VPN	Impacted	Connectivity between VPN peers (peer VRFs hosted by PE routers) in this VPN is impaired by a connectivity failure in the transport layer.
VRF	Impacted	Connectivity between this VRF and one or more of its VPN peers is impaired by a connectivity failure in the transport layer.

L2VPN Domain Impact Events

Table 22 lists the impact events notified by the MPLS Manager for the L2VPN domain, including the condition for each event. The table also identifies the managed elements for which the impact events are notified.

Table 22: L2VPN Impact Events Notified by the MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
Forwarder	Impacted	This Forwarder is impaired by: <ul style="list-style-type: none"> • The failure of the PE router hosting this Forwarder or • The failure of the PE router hosting the peer Forwarder.
PseudoWire	Impacted	Connectivity for this Pseudowire is impaired by: <ul style="list-style-type: none"> • The failure of at least one of the corresponding terminating Forwarders, • The failure of an underlying LSP, or • The failure of an underlying LdpAdjacency.
LdpProtocolEndpoint	Impacted	This LdpProtocolEndpoint is impaired by an underlying physical failure.

Table 22: L2VPN Impact Events Notified by the MPLS Manager *(continued)*

MANAGED ELEMENT	EVENT	CONDITION
VLAN	Impacted	This VLAN is impaired by an underlying: <ul style="list-style-type: none"> • PseudoWire Down event, • PseudoWire Impacted event, • Forwarder IsDown event, or • Forwarder Impacted event.

MPLS Domain Impact Events

Table 23 lists the impact events notified by the MPLS Manager for the MPLS domain, including the condition for each event. The table also identifies the managed elements for which the impact events are notified.

Table 23: MPLS Impact Events Notified by the MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
LSP	Impacted	Connectivity for this path is impaired by a connectivity failure in the transport layer: The two PE routers serviced by this LSP can no longer communicate via this LSP.

Impact Analysis Examples for the L3VPN Domain

The two examples that follow show how the MPLS Manager uses the relationships between the underlying transport network elements and the VPN, VRF, and LSP elements to perform MPLS impact analysis. In addition, the examples show how the Global Manager picks up where the MPLS Manager leaves off to perform global impact analysis.

Figure 12 shows the underlying transport failures for the two examples.

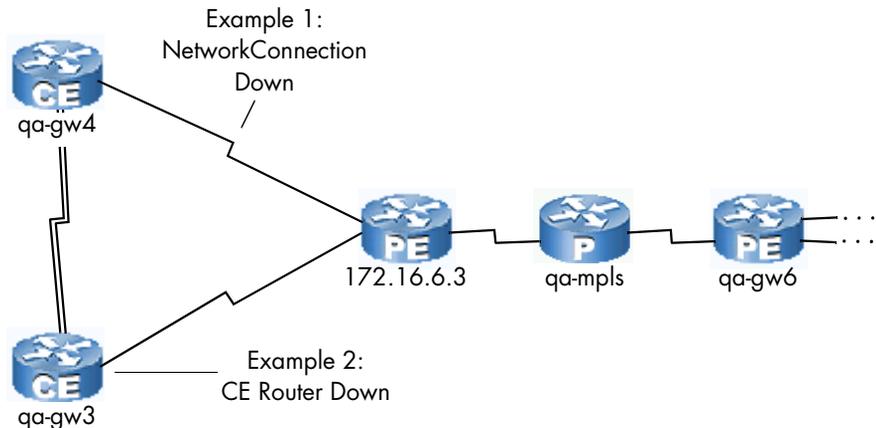


Figure 12: Underlying Transport Failures for the L3VPN Impact Examples

L3VPN Impact Example 1: CE-PE Router NetworkConnection Down

In this example, the MPLS Manager receives a NetworkConnection Down problem from the Availability Manager for a failed connection between a CE router and a PE router.

When the Availability Manager detects the NetworkConnection Down problem, it sends the problem along with the network symptoms to both the Global Manager and the MPLS Manager. The MPLS Manager calculates the MPLS impacts caused by the problem and sends an impact notification for each of the impacts to the Global Manager. Each impact notification lists the NetworkConnection Down problem as the cause of the impact.

The Global Manager adds the impacts in the impact notifications received from the MPLS Manager to the symptoms in the root-cause problem notification received from the Availability Manager to form a combined list of impacts for the root-cause problem notification.

What follows is a summary of example notifications created for the NetworkConnection Down problem, followed by an example display (Figure 13) showing the combined impacts for the NetworkConnection Down problem notification.

- Availability Manager Root-Cause Notification:**
 Root Cause: NetworkConnection Down
 Symptom: NetworkConnection DownOrFlapping
- MPLS Manager Impact Notifications:**
 Impact: VPN Impacted
 Impact: VRF Impacted
- Global Manager Modified Root-Cause Notification:**
 Root Cause: NetworkConnection Down
 Symptom: NetworkConnection DownOrFlapping
 Impacts: VPN Impacted and VRF Impacted

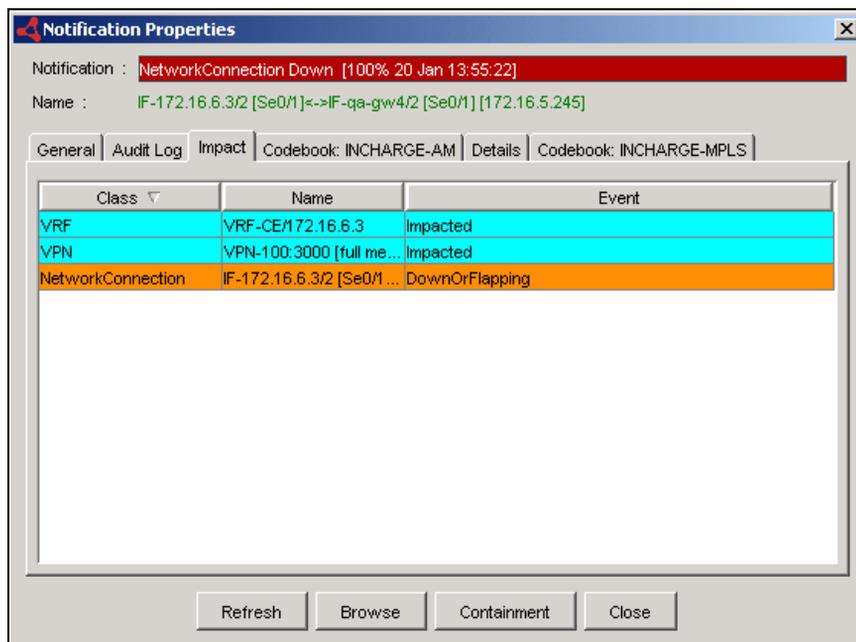


Figure 13: Notification Properties Dialog Box Showing a NetworkConnection Down Problem

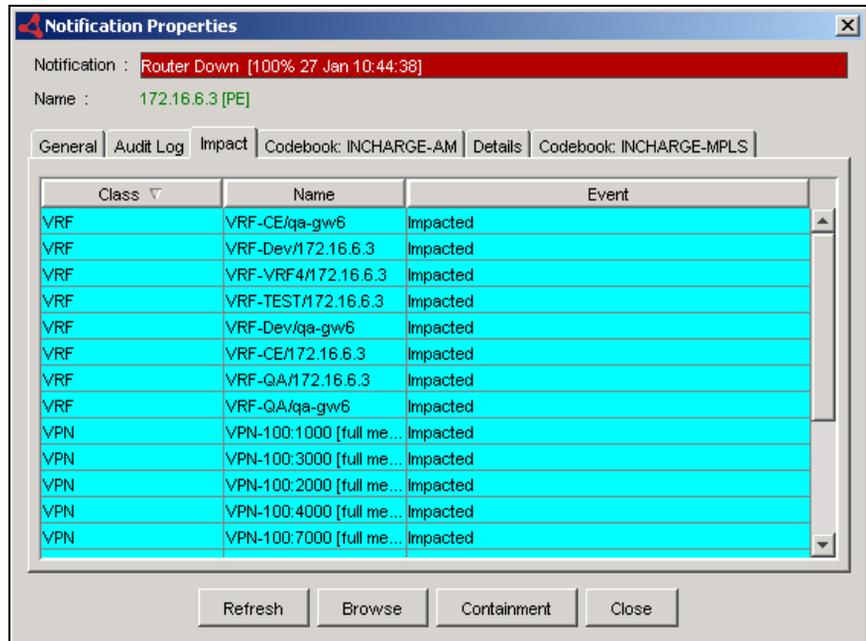
For instructions on viewing detailed notification information, see [Viewing MPLS Notifications, Maps, and Containment](#) on page 49.

L3VPN Impact Example 2: PE Router Down

In this example, the MPLS Manager receives a Router Down problem from the Availability Manager for a failed PE router.

What follows is a summary of example notifications created for the Router Down problem, followed by an example display (Figure 14) showing the combined impacts for the Router Down problem notification.

- **Availability Manager Root-Cause Notification:**
Root Cause: Router Down
Symptoms: NetworkConnection DownOrFlapping, one or more instances of Router Unresponsive
- **MPLS Manager Impact Notifications:**
Impact: One or more instances of LSP Impacted
Impact: One or more instances of VPN Impacted
Impact: One or more instances of VRF Impacted
- **Global Manager Modified Root-Cause Notification:**
Root Cause: Router Down
Symptoms: NetworkConnection DownOrFlapping and one or more instances of Router Unresponsive
Impacts: One or more instances of LSP Impacted, VPN Impacted, and VRF Impacted



Continue

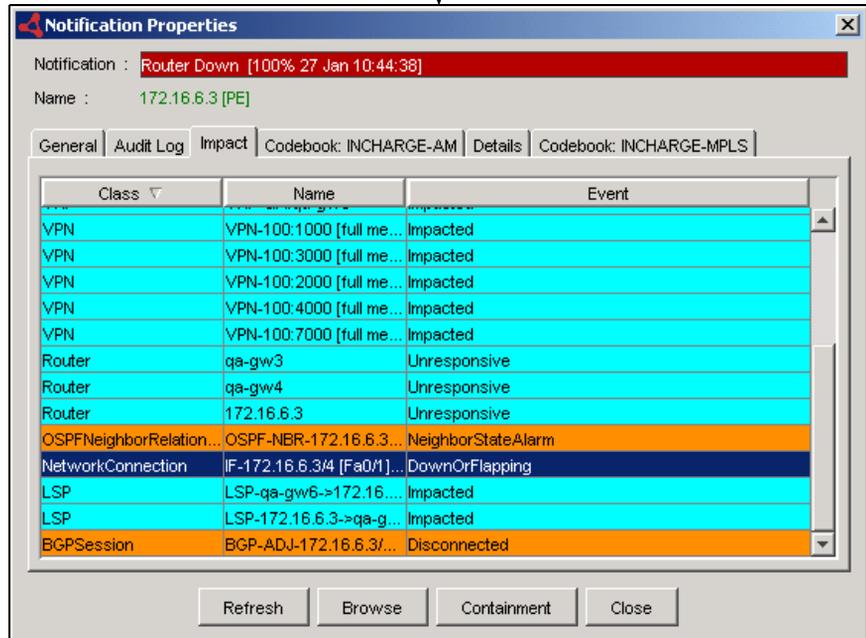


Figure 14: Notification Properties Dialog Box Showing a Router Down Problem

Impact Analysis Examples for the L2VPN Domain

The following example shows how the MPLS Manager uses the relationships between the Forwarders and PseudoWires with the underlying VLAN to perform impact analysis using the Transparent LAN Service (TLS) feature. In this example, when an interface on a PE router goes down, the Forwarder's relationships with the PseudoWire and the associated VLAN are impacted.

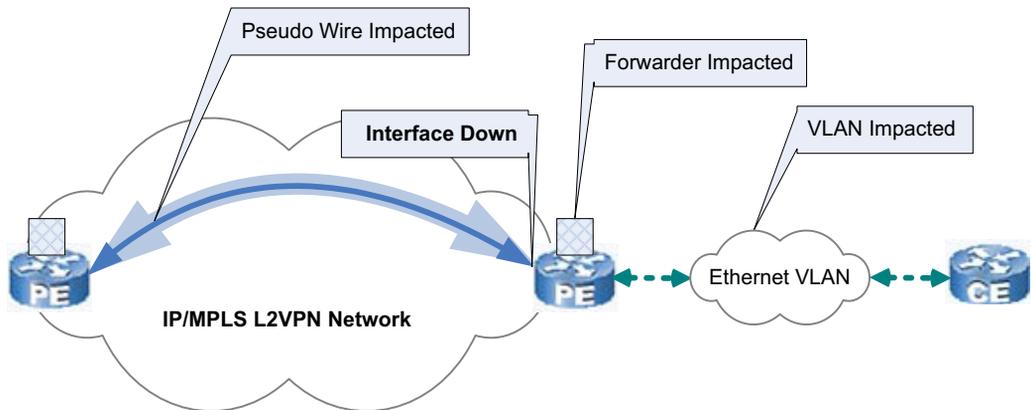


Figure 15: Example Showing How Interface Down Impacts L2VPN Elements

In this example, the MPLS Manager receives an Interface Down problem from the Availability Manager for a failed interface on a PE router.

When the Availability Manager detects the Interface Down problem, it sends the problem along with the network symptoms to both the Global Manager and the MPLS Manager. The MPLS Manager calculates the L2VPN impacts caused by the problem and sends an impact notification for each of the impacts to the Global Manager. Each impact notification lists the Interface Down problem as the cause of the impact.

The Detail tab for both the Forwarder Impacted notification and the PseudoWire Impacted notification indicates the VLANs associated with the Forwarder Impacted events. Because this impact event is caused by a physical failure originating in the Availability Manager domain, the Caused By tab will indicate the source of the event; in this example, Interface Down.

What follows is a summary of example notifications created for an Interface Down problem.

- **Availability Manager Root-Cause Notification:**
Root Cause: Interface Down
Symptom: Interface DownOrFlapping
- **MPLS Manager Impact Notifications:**
Impact: VLAN Impacted
Impact: Forwarder Impacted
Impact: PseudoWire Impacted
- **Global Manager Modified Root-Cause Notification:**
Root Cause: Interface Down
Symptom: Interface DownOrFlapping
Symptoms: NetworkConnection DownOrFlapping and one or more instances of Router Unresponsive
Impacts: One or more instances of LSP Impacted, VPN Impacted, and VRF Impacted

For instructions on viewing detailed notification information, see [Viewing MPLS Notifications, Maps, and Containment](#) on page 49.

4

Viewing MPLS Notifications, Maps, and Containment

This chapter describes using the Global Console to view MPLS notifications, topology maps, and containment information for the MPLS Manager. The MPLS topology maps include:

- LSP maps
- LSP Hop maps
- VPN maps
- PseudoWire maps

You view the notifications, maps, and containment information by attaching the Global Console to the Global Manager.

For instructions on viewing notifications, maps, and containment information, see the *EMC Smarts Service Assurance Manager Operator's Guide*. For information about the topology elements discovered and the events notified by MPLS Manager, see [MPLS and VPN Elements and Their Failures](#) on page 9.

Viewing MPLS Notifications

The MPLS Manager reports notifications to the Global Manager, and the Global Manager combines these notifications with the notifications received from the Availability Manager. You can view the notifications through the Global Console in two basic ways:

- As table entries in a Notification Log Console view
- As color-coded severity icons in a Map Console view

Opening an MPLS Notification Properties Dialog Box

To obtain detailed information about an individual MPLS notification, you can use any of the following common methods to open the Notification Properties dialog box:

- Double-click an MPLS notification in the Notification Log Console.
- Select an MPLS notification in the Notification Log Console and click the **Properties** toolbar button.
- Right-click a selected MPLS notification and select *Properties* in the pop-up menu.
- Double-click an MPLS map icon affected by active events.

MPLS Notification Properties

Figure 16 and Figure 17 provide examples of detailed VRF Down and VRF NoRoutes information displayed in the Notification Properties dialog box.

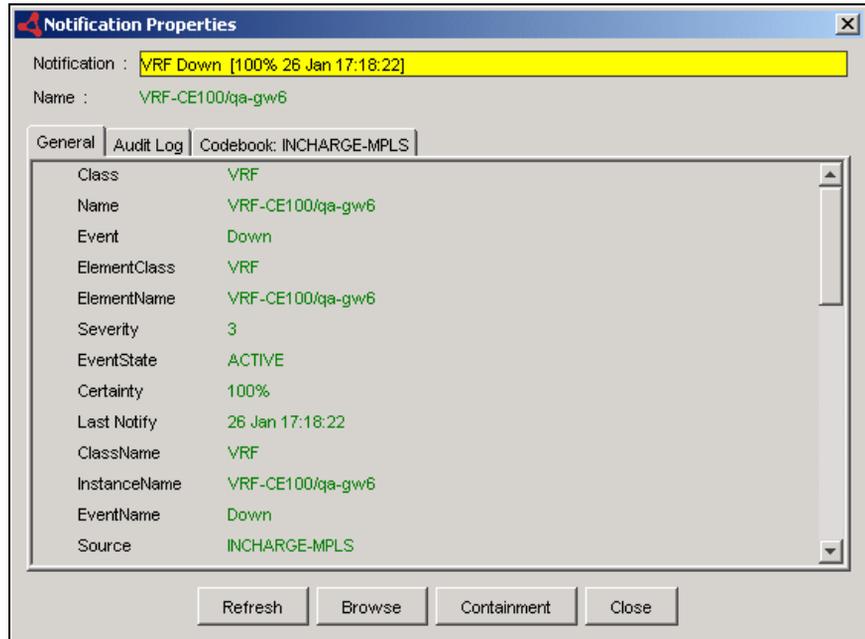


Figure 16: Notification Properties Dialog Box Showing a Down VRF

The VRF in Figure 16 either has no associated interfaces or has one or more associated interfaces and all of them are down.

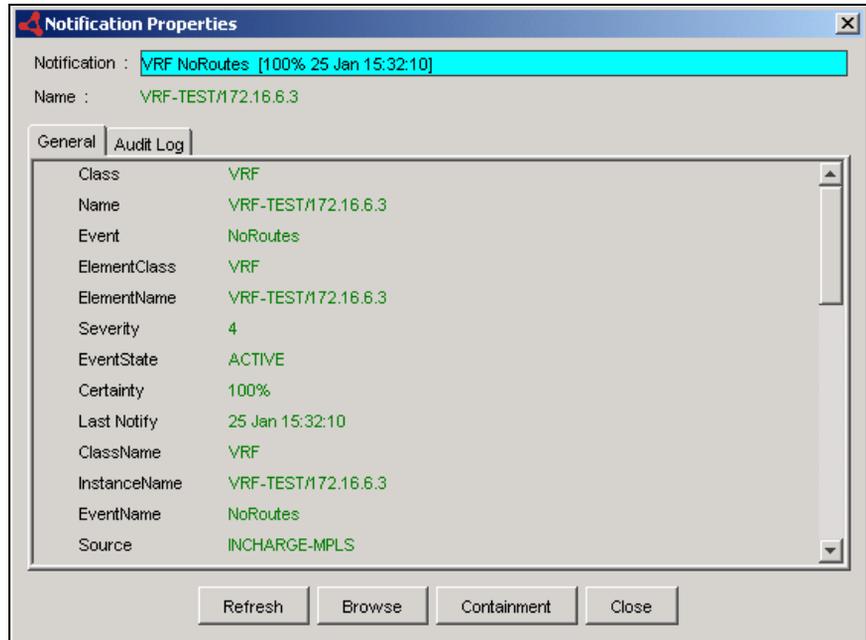


Figure 17: Notification Properties Dialog Box Showing a Routeless VRF

The VRF in Figure 17 has one or more associated interfaces but all of them are unnumbered. (An unnumbered interface has no IP address assigned to it.) For a VRF to be up (operational), it must have at least one numbered interface that is up.

Viewing MPLS Topology in Maps

The MPLS Manager sends a streamlined copy of its MPLS topology—instances of the MPLSService, LSP, LSPHop, VPN, VRF, Forwarder, and PseudoWire classes—to the Global Manager. The Global Manager combines this topology with the underlying transport network topology received from the Availability Manager.

The Global Console presents the topology information in a variety of dynamically updated formats that show the status of the MPLS and network elements and their many relationships. One of those formats is the topology map, which is a graphical representation of the topology.

Viewing topology maps is an easy and quick way to learn more about the source, impact, and cause of MPLS notifications. You view the MPLS topology maps using the Map Console view of the Global Console.

Opening an MPLS Topology Map

You can use any of the following common methods to open an MPLS topology map:

- Open the Map Console by selecting the Show Map option from any opened console attached to the Global Manager. For example, in the Notification Log Console, click an MPLS notification and then select *Event > Show Map*, or right-click the notification and then select *Show Map* in the pop-up menu. In the Topology Browser Console, right-click an MPLS element (router, LSP, VPN, VRF, Forwarder, or PseudoWire) and select *Show Map* in the pop-up menu.
- Open the Map Console from the Global Console by selecting *File > New > Map Console*. In the Topology tab of the Map Console, click an MPLS element (router, LSP, VPN, VRF, Forwarder, or PseudoWire) to display a map for the element, or right-click an MPLS element (router, LSP, VPN, VRF, Forwarder, or PseudoWire) and select an MPLS map type (LSP, LSP Hops, VPN, or PseudoWire) from the pop-up menu.
- In an opened topology map, right-click an MPLS map icon and select an MPLS map type (LSP, LSP Hops, VPN, or PseudoWire) from the pop-up menu.

MPLS Topology Map Graphical Representations

MPLS topology maps contain router, LSP, LSP hop, VPN, VRF, Forwarder, and PseudoWire elements, along with their relationships and connections. In a map display, a *node* is a graphical representation of an element, and an *edge* is a graphical representation of a relationship or connection between elements.

Table 24 identifies and describes the default nodes and edges that may appear in an MPLS topology map. In the Map Console, you can also select *Map > Map Legend* to see a similar list.

Note that your system administrator may replace the standard map nodes with other map nodes that are preferred by your organization. In that case, use *Map > Map Legend* to see the definitions of your map nodes.

Table 24: Default Nodes and Edges for MPLS Topology Maps

ICON / VISUAL INDICATOR	DESCRIPTION
	<p>Standard router node with PE inscription—represents a Provider Edge (PE) router and the MPLS service element associated with the PE router.</p>
	<p>Standard router node with P inscription—represents a Provider (P) router and the MPLS service element associated with the P router.</p>
	<p>Standard router node with CE inscription—represents a Customer Edge (CE) router and the MPLS service element associated with the CE router.</p>
	<p>Represents a Layer 3 (L3) or Layer 2 (L2) VPN.</p>
	<p>Represents a VRF in an L3VPN-related map.</p>
	<p>Represents a Forwarder in an L2VPN-related map.</p>
	<p>Solid line can represent a physical connection, a logical IP connection, a logical VLAN connection, a membership, or a group relationship.</p>
	<p>Jagged line can represent a network connection between routers or a virtual link between a VRF and a CE router.</p>

Table 24: Default Nodes and Edges for MPLS Topology Maps (*continued*)

ICON / VISUAL INDICATOR	DESCRIPTION
	Solid black line with arrow can represent a dependency or an LSP in <i>No Highlight LSP</i> mode on an LSP Hops map.
	Dotted black line with arrow can represent composition or an LSP in <i>Highlight LSP</i> mode. When representing an LSP in <i>Highlight LSP</i> mode on the LSP Hops map, the line animates to show the direction of packet flow through the LSP.

Note: Additional icons may display, depending on the underlying SMARTS products and certified devices.

MPLS Map Types

Maps for MPLS show the routers and LSPs that support the VPNs, VRFs, Forwarders, and PseudoWires. There are four types of maps that focus on MPLS elements:

- LSP map—available to all MPLS elements (see Note); default map for Forwarders
- LSP Hops map—available to all MPLS elements; default map for LSPs
- VPN map—not available to LSPs; default map for VPNs and VRFs
- PseudoWire map—not available to LSPs or VRFs; default map for PseudoWires

Note: No MPLS map types are available to MPLSService and LSPHop.

LSP Map

The LSP map (Figure 18) shows the LSP connectivity between PE routers.

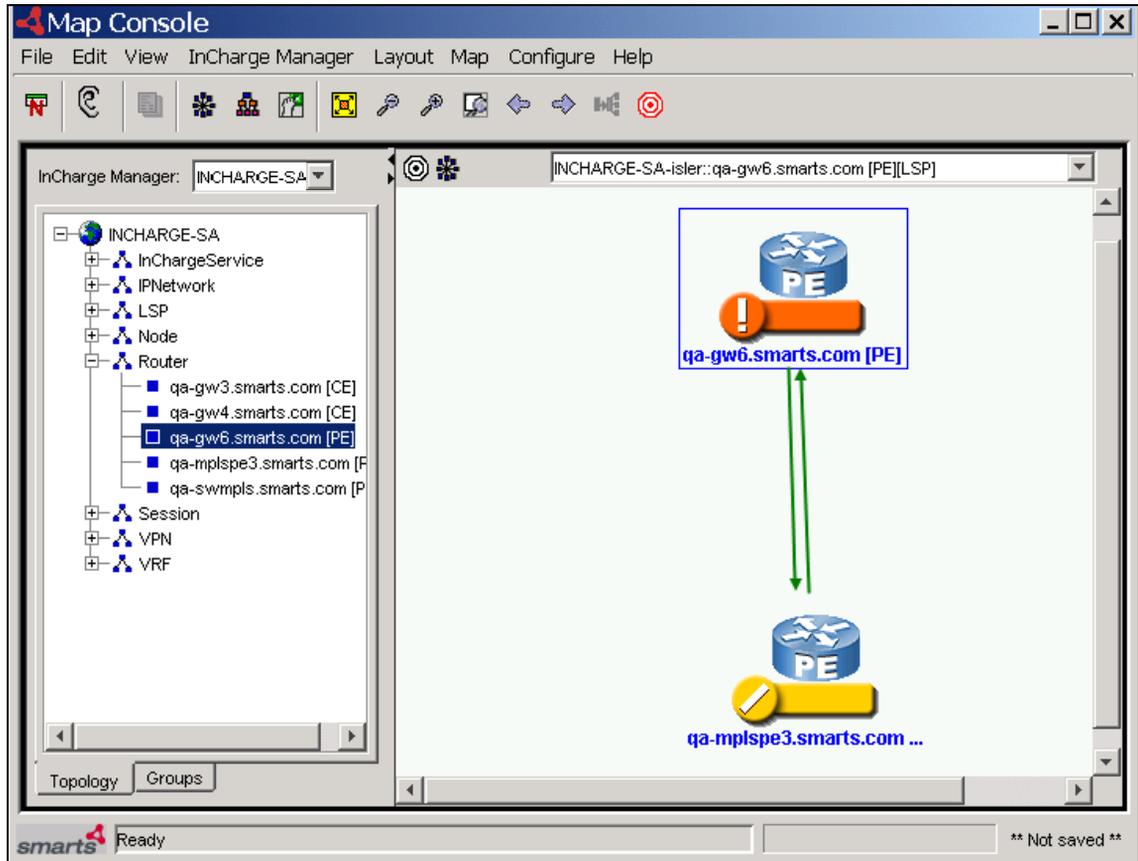


Figure 18: LSP Map Launched from a PE Router Element—Example

An arrowhead line in the LSP map represents an LSP instance; the arrowhead signifies the destination end of the LSP. Because an LSP instance is from PE router to PE router, no P routers appear in the LSP map.

As clarified in Table 25, the actual display of an LSP map depends on the source element from which the map is launched.

Table 25: LSP Map Display When Launched from Different Source Elements

WHEN LAUNCHED FROM THIS ELEMENT . . .	THE MAP DISPLAYS . . .
PE router	All of the LSPs that either originate or terminate at the PE router.
LSP	The LSP and the two PE routers associated with it.
VPN	The PE routers and the LSPs belonging to the VPN in either the full-mesh configuration or hub-and-spoke configuration.
VRF	In an L3VPN, all of the LSPs used by the VRF to communicate with its peer VRFs.
Forwarder	In an L2VPN, all of the LSPs used by the Forwarder to communicate with its peer Forwarder.
PseudoWire	All of the LSPs underlying the PseudoWire.

LSP Hops Map

The LSP Hops map (Figure 19) shows the intermediate LSP hops that comprise an LSP.

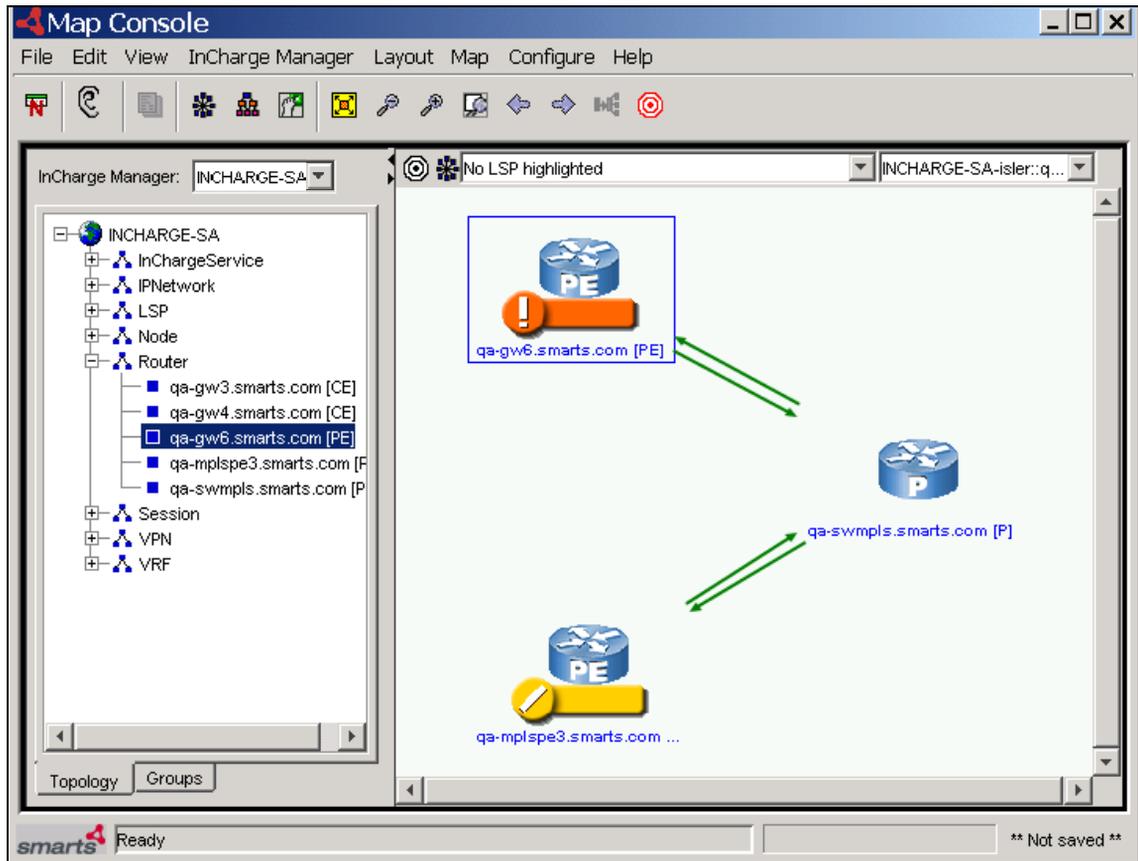


Figure 19: LSP Hops Map Launched from a PE Router Element—Example

The LSP Hops map displays information similar to that of the LSP map except that instead of displaying just the endpoint PE routers of an LSP, it includes all the transit P routers as well. And, as with the LSP map, the actual display of an LSP Hops map depends on the source element from which the map is launched.

At the top of an LSP Hops map is a Highlight LSP drop-down box containing the names of the LSPs appearing in the map. By default, the *No LSP highlighted* option is selected, and the arrowhead lines representing the LSP hops are solid lines. Selecting an LSP name from this drop-down box causes all the arrowhead lines representing the hops for that LSP to change to animated dotted lines, to show the flow of packets through the LSP.

VPN Map for an L3VPN Network

The VPN map for an L3VPN network (Figure 20) shows the VPN, the VRFs that are members of the VPN, the PE routers that host the VRFs, and the CE routers to which the VRFs virtually connect.

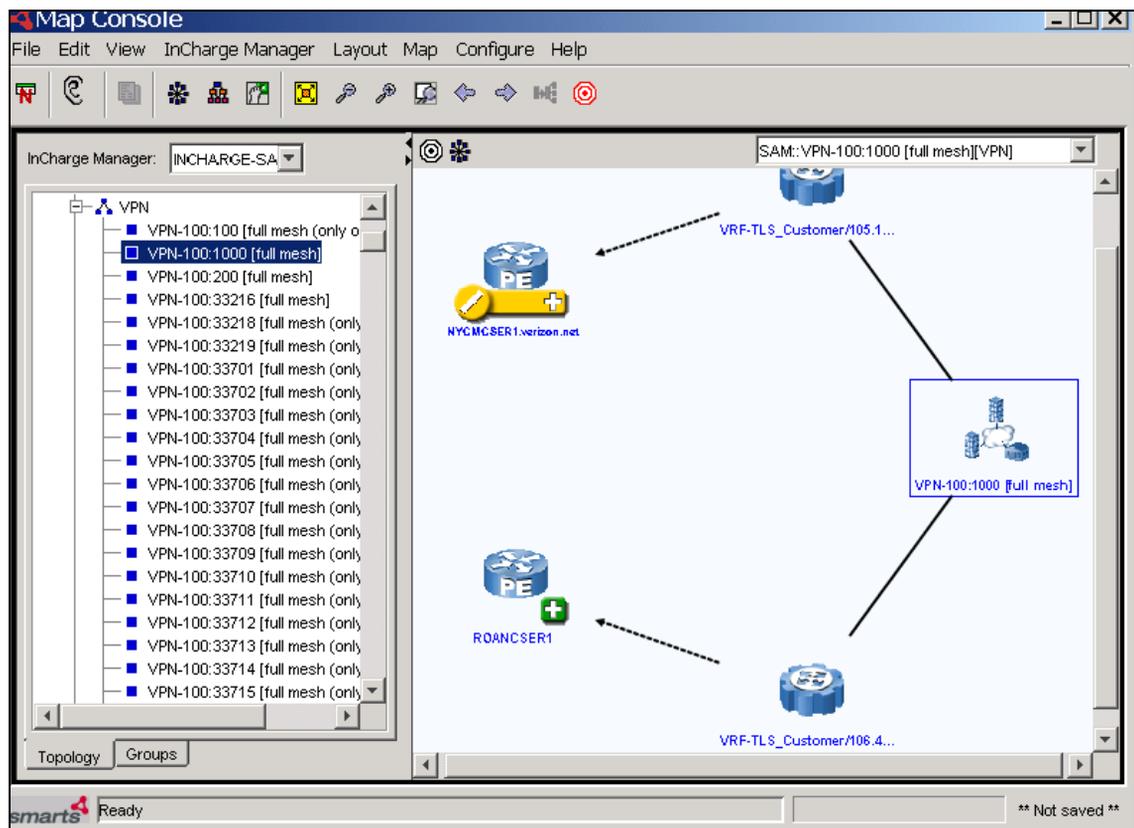


Figure 20: VPN Map Launched from a Layer 3 VPN Element—Example

As clarified in Table 26, the actual display of a VPN map for an L3VPN network depends on the source element from which the map is launched.

Table 26: VPN Map Display for L3VPN When Launched from Different Source Elements

WHEN LAUNCHED FROM THIS ELEMENT . . .	THE MAP DISPLAYS . . .
PE router	All the VRFs hosted by the PE router.
CE router	The VRF to which the CE router is virtually connected.
VPN	All the VRFs that are part of the VPN and the PE routers hosting the VRFs.
VRF	The VRF, the PE router hosting the VRF, and the CE routers to which the VRF virtually connects.

If additional VRFs (that belong to other VPNs) are hosted by the PE routers, the map can be expanded to see those VRFs as well.

VPN Map for an L2VPN Network

The VPN map for an L2VPN network (Figure 21) shows the VPN, the Forwarders and PseudoWires that are members of the VPN, the PE routers that host the Forwarders, and the CE routers to which the Forwarders virtually connect.

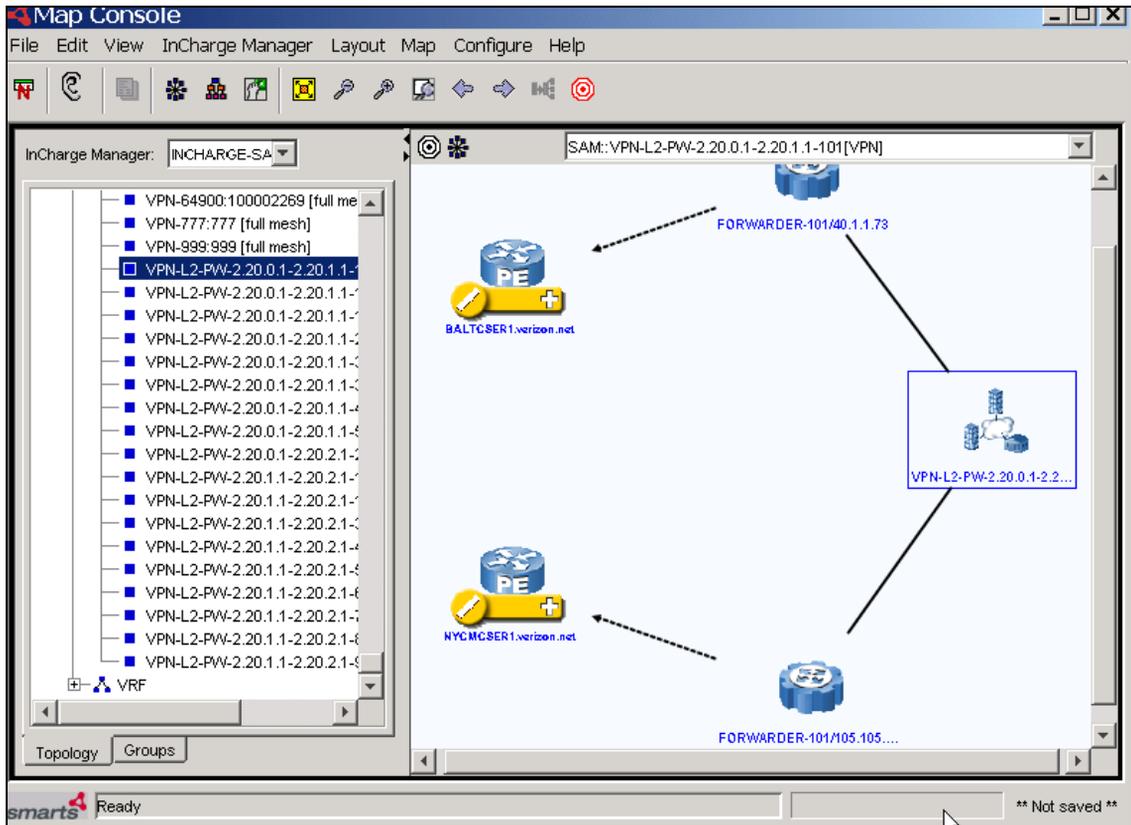


Figure 21: VPN Map Launched from a Layer 2 VPN Element—Example

As clarified in Table 27, the actual display of a VPN map for an L2VPN network depends on the source element from which the map is launched.

Table 27: VPN Map Display for L2VPN When Launched from Different Source Elements

WHEN LAUNCHED FROM THIS ELEMENT . . .	THE MAP DISPLAYS . . .
PE router	All the Forwarders hosted by the PE router.
CE router	The Forwarder to which the CE router is virtually connected.
VPN	All the Forwarders and PseudoWires that are part of the VPN and the PE routers hosting the Forwarders.
Forwarder	The Forwarder, the PE router hosting the Forwarder, and the CE routers to which the Forwarder virtually connects.

If additional Forwarders (that belong to other VPNs) are hosted by the PE routers, the map can be expanded to see those Forwarders as well.

VPN Map PE/CE Display Toggle Feature

You can use the *Show PEs* and *Show CEs* menu options (Figure 22) to temporarily hide all PEs and/or CEs in a VPN map display.

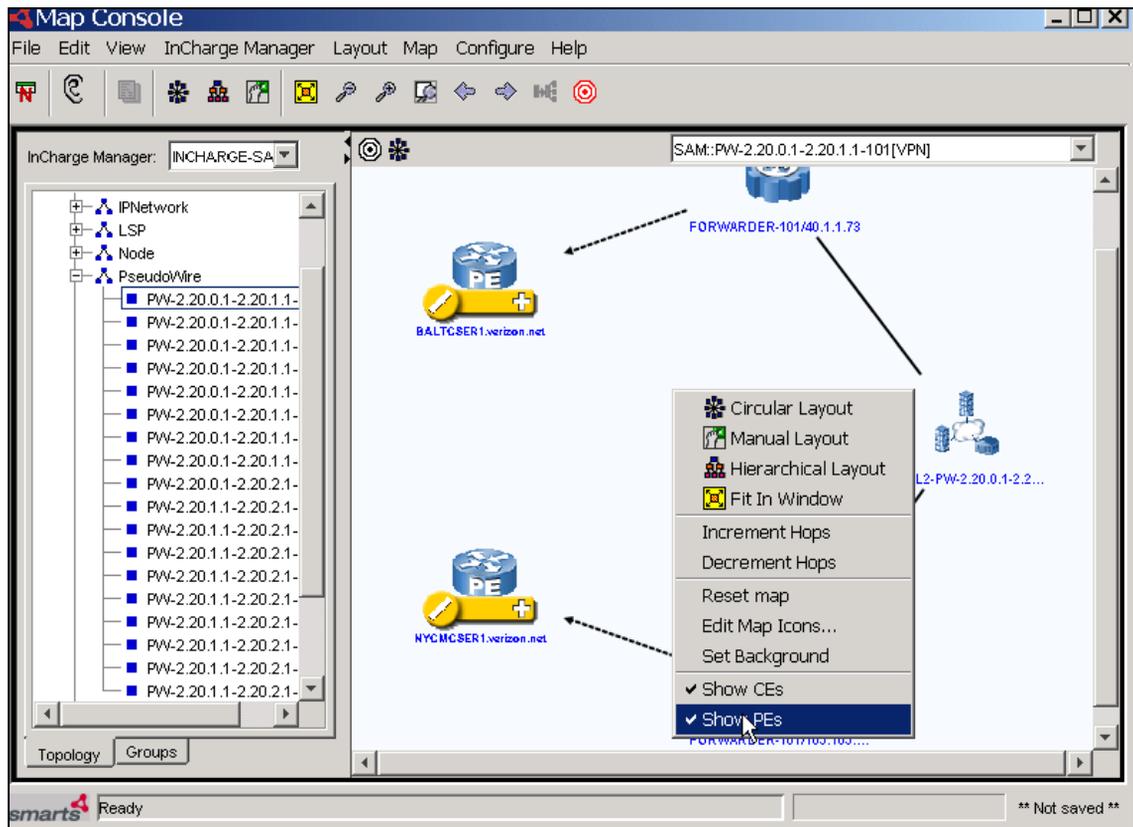


Figure 22: VPN Map Showing PE/CE Toggle Option

To hide the PEs or CEs in the VPN map display, right-click an open space in the map display and select the *Show PEs* or *Show CEs* option from the pop-up menu. To show the PEs or CEs, right-click an open space in the map display and select the *Show PEs* or *Show CEs* option again.

PseudoWire Map

The PseudoWire map (Figure 23), applicable only to L2VPN networks, shows solid lines with arrows to represent the LSPs underlying the PseudoWires.

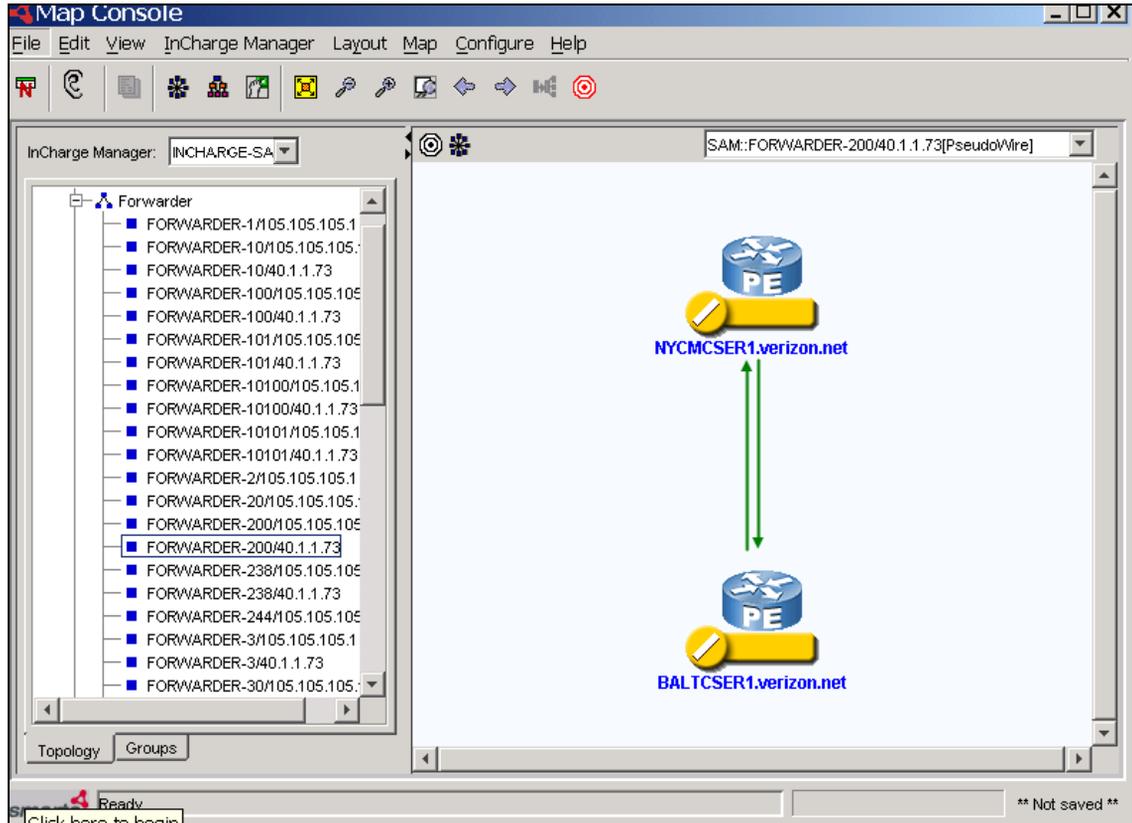


Figure 23: PseudoWire Map Launched from a Forwarder—Example

As clarified in Table 28, the actual display of a PseudoWire map depends on the source element from which the map is launched.

Table 28: PseudoWire Map Display When Launched from Different Source Elements

WHEN LAUNCHED FROM THIS ELEMENT . . .	THE MAP DISPLAYS . . .
PE router	The LSPs underlying the PseudoWires that terminate on this PE router.
VPN	The LSPs underlying the PseudoWire associated with this VPN.
Forwarder	The LSPs underlying the PseudoWire associated with this Forwarder.

Enhanced VPN Maps

Enhanced VPN maps supporting customer business service views are available to MPLS Manager deployments that include the Business Impact Manager and a specialized EMC Smarts adapter that interfaces with the customer's provisioning system. Currently, only one such adapter is available: Adapter for Cisco ISC, which interfaces with the Cisco Internet Solutions Center (ISC) provisioning system.

For information about the enhanced VPN maps available through the Adapter for Cisco ISC, see the *EMC Smarts Adapter for Cisco ISC User's Guide*. For information about the Business Impact Manager, see the *EMC Smarts Business Impact Manager User's Guide*.

Viewing MPLS Containment

Accompanying the MPLS topology that the MPLS Manager sends to the Global Manager are instructions on what topology information to include in the Containment view for each type of MPLS element. An element's Containment view displays in a dialog box that organizes the information for the element into tab pages and tables. Figure 24 shows a Containment dialog box for an LSP element.

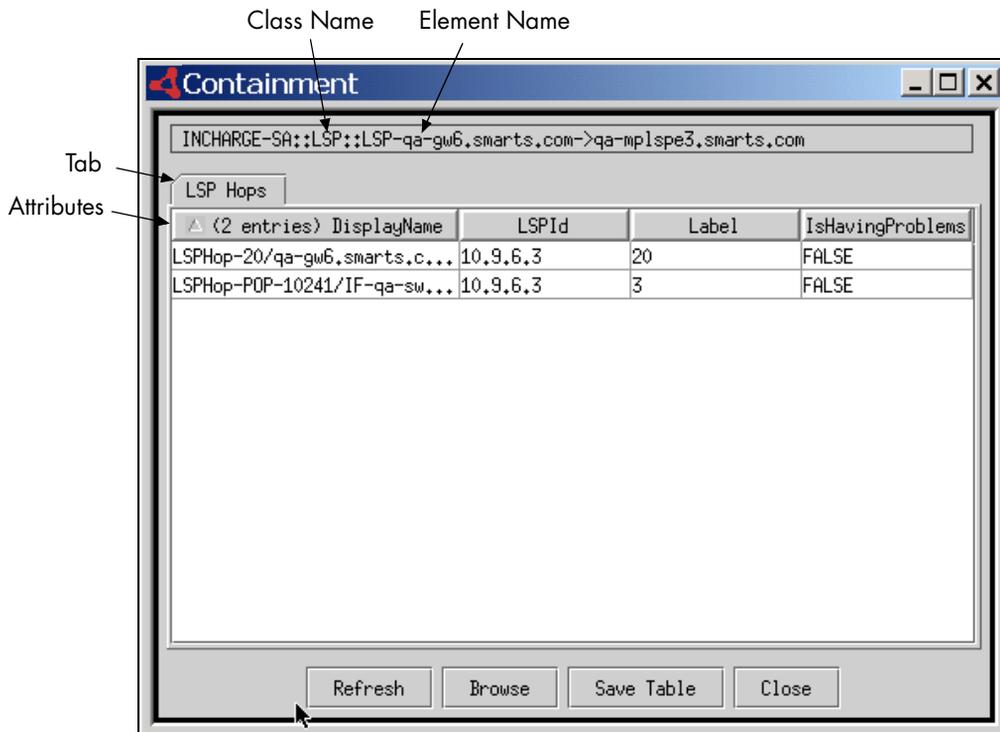


Figure 24: LSP Containment Dialog—Example

In a Containment view, different information is available for different types of elements.

Note: Containment is not available for user-defined groups and services.

Opening an MPLS Containment Dialog

To obtain containment information about an individual MPLS element, you can use any of the following common methods to open a Containment dialog for the element:

- In an MPLS map tree, right-click the element and select Containment in the pop-up menu.
- In an MPLS map display, right-click the element and select Containment in the pop-up menu.

- In a Topology Browser Console, right-click the element and select Containment in the pop-up menu.
- In a Notification Log Console, double-click a notification for the element and click the **Containment** button in the Notification Properties dialog.

MPLS Containment Tab Pages

The MPLS tab pages that appear in a Containment view for an MPLS element contain the routing topology information listed in the following table.

Table 29: MPLS Classes and Their Containment Tab Pages and Associated Attributes

CLASS NAME	MPLS TAB NAME	ATTRIBUTES
LSP	LSP Hops	DisplayName LSPId Label IsHavingProblems
VPN	VPN PEs	DisplayName Description Vendor Model Type Location
	All VRFs	VRFName SystemName DisplayName RouteDistinguisher OperStatus NumberOfRoutes

Table 29: MPLS Classes and Their Containment Tab Pages and Associated Attributes *(continued)*

CLASS NAME	MPLS TAB NAME	ATTRIBUTES
VRF	VPN Peers	VRFName SystemName DisplayName RouteDistinguisher OperStatus NumberOfRoutes
	VRF Interfaces	CreationClassName Type DisplayName AdminStatus
	CEs Attached	VRF CEs ICIM_UnitaryComputerSystem DisplayName Description Vendor Model Type Location

Customizing Groups and Settings

The MPLS Manager monitors the MPLS network by (1) polling the SNMP agent on the discovered routing devices for Forwarder, LdpProtocolEndpoint, and VRF information and then (2) comparing the polling results to threshold values that define acceptable and unacceptable levels of Forwarder, LdpProtocolEndpoint, and VRF operation. The MPLS Manager uses the thresholds values, in conjunction with events from Availability Manager, to monitor Forwarder, LdpProtocolEndpoint, and VRF availability and to detect VRF space overload.

Also, if periodic remote pings are deployed (see Figure 25), the MPLS Manager monitors reachability between the devices in the MPLS network by sending remote ping requests to the discovered PE, P, and CE devices, polling the ping test results, and comparing the ping test results to threshold values that define acceptable and unacceptable levels of reachability. MPLS Manager uses the thresholds values, in conjunction with events from Availability Manager, to monitor PE, P, and CE interconnectivity.

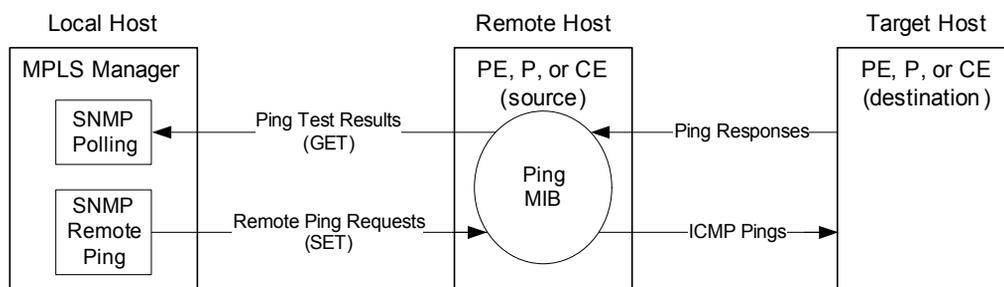


Figure 25: MPLS Manager Remote Ping Creation and Monitoring

Using the Polling and Thresholds Console, shown in the figure [Polling and Thresholds Console—Example](#) on page 89, you can customize the SNMP polling for MPLS Manager by modifying default polling groups or by creating new polling groups. The polling groups periodically poll the managed MPLS environment to collect the data needed by MPLS Manager to determine the availability of the Forwarder, LdpProtocolEndpoint, and VRF elements. The data serves as input to MPLS Manager correlation analysis.

Using the Polling and Thresholds Console, you can also deploy and customize periodic remote pings, and customize the SNMP polling of the remote ping test results, by modifying default remote ping groups or by creating new remote ping groups. The remote ping groups trigger ping tests on PE, P, and CE devices in the managed MPLS environment and periodically poll the managed MPLS environment to collect the ping test results needed by MPLS Manager to determine the interconnectivity of the PE, P, and CE elements. The ping test results serve as input to MPLS Manager correlation analysis.

For a description of remote ping operation, see [Remote Ping Functionality](#) on page 97. For a description of remote ping enabling and global-value customizing, see the *EMC Smarts MPLS Manager Configuration Guide*. For a description of SNMP polling for correlation analysis, see [Polling for Analysis](#) on page 127.

Default Groups and Settings

The polling and remote ping groups and settings for the MPLS Manager are accessible via the Polling tab of the Polling and Thresholds Console. A *group* contains one or more settings, and a *setting* contains a collection of parameters.

The MPLS Manager provides nine group categories, shown in Figure 26, each of which contains a default group.

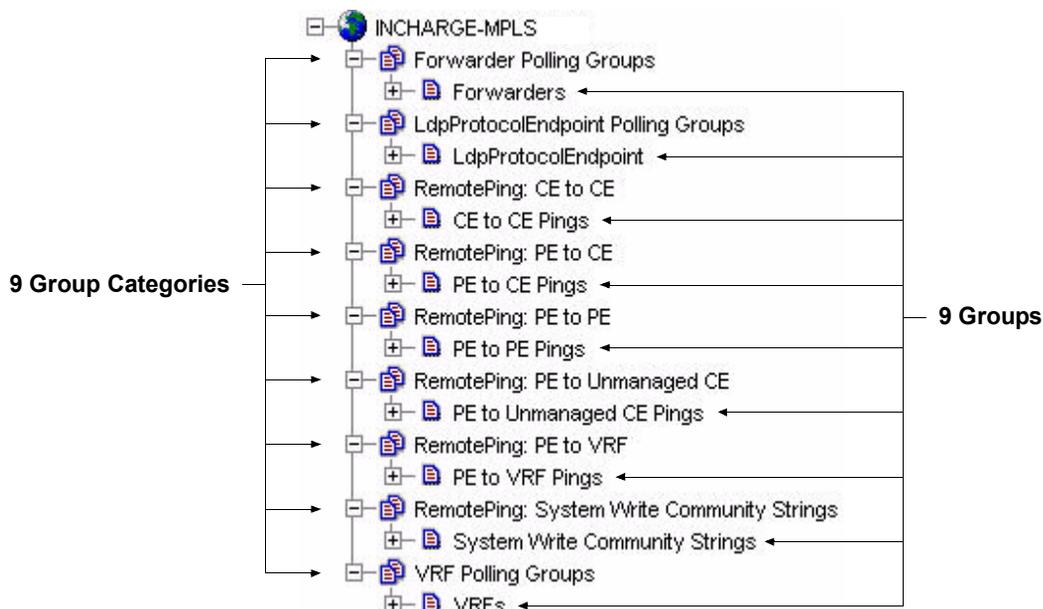


Figure 26: MPLS Manager Group Categories and Default Groups

Table 30 lists the default polling and remote ping groups and their settings.

Table 30: Default Groups and Settings

GROUPS	TARGET CLASS	SETTINGS
<i>Default Polling Groups</i>		
Forwarder	Forwarder	Forwarder SNMP Setting
LdpProtocolEndpoint	LdpProtocolEndpoint	LDP Session SNMP Setting
VRF	VRF	VRF SNMP Setting (default setting) VRF External Setting
<i>Default Remote Ping Groups</i>		
CE to CE	VPN	CE to CE Ping Setting
PE to CE	VPN	PE to CE Ping Setting
PE to PE	VPN	PE to PE Ping Setting
PE to Unmanaged CE	VPN	PE to Unmanaged CE Ping Setting

Table 30: Default Groups and Settings (*continued*)

GROUPS	TARGET CLASS	SETTINGS
PE to VRF	VPN	PE to VRF Ping Setting
System Write Community Strings	Unitary Computer System	System Writer Community Strings

Three things to note about this table:

- Even though all five “Ping” settings (*CE to CE Ping Setting* through *PE to VRF Ping Setting*) are available to each of the default remote ping groups (excluding the *System Write Community Strings* group), only the group-to-setting mappings shown in the table are valid.
- Only the discovered elements of the type identified in the *Matching Criteria* column for a particular group can become members of that group. Membership for type *VPN* is limited to layer 3 VPNs (L3VPNs).
- By default, *all* discovered elements of the type identified in the *Matching Criteria* column for a particular group become members of that group.

You can specify additional matching criteria for a polling or remote ping group to further limit its membership, as explained in [Editing Matching Criteria](#) on page 92. For a remote ping group, you can also configure certain setting parameters to limit its membership.

Forwarders

The *Forwarders* default polling group is used by the MPLS Manager to monitor the availability of Forwarder instances discovered on Juniper routing devices. It has a single setting named *Forwarder SNMP Setting*.

This setting determines the polling intervals used by the MPLS Manager to monitor Forwarders. MPLS Manager monitors the Forwarders by probing the SNMP tables on the PE routers hosting the Forwarders.

Table 31 lists the *Forwarder SNMP Setting* parameters.

Table 31: Forwarder SNMP Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: ENABLED	Enables or disables the MPLS Manager polling of Forwarders and the subsequent analysis of polled Forwarder data.
PollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls.
Retries	0 to 10 retries Default: 3	Sets the number of retry polls to perform when the initial poll fails.
Timeout	10 to 10000 milliseconds Default: 700 milliseconds	Sets the amount of time to wait for the poll response before the first poll request times out. The timeout value doubles for each successive retry. For Timeout=700 msec (0.7 sec) and Retries=3: <ul style="list-style-type: none"> • 0.7 seconds for first retry • 1.4 seconds for second retry • 2.8 seconds for third retry

LdpProtocolEndpoint

The *LdpProtocolEndpoint* default polling group is used by the MPLS Manager to monitor the availability of LdpProtocolEndpoint instances discovered on Cisco routing devices. It has a single setting named *LDP Session SNMP Setting*.

This setting determines the polling intervals used by the MPLS Manager to monitor LdpProtocolEndpoints. MPLS Manager monitors the LdpProtocolEndpoints by probing the SNMP tables on the PE routers hosting the LdpProtocolEndpoints.

Table 32 lists the *LDP Session SNMP Setting* parameters.

Table 32: LDP Session SNMP Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: ENABLED	Enables or disables the MPLS Manager polling of LdpProtocolEndpoints and the subsequent analysis of polled LdpProtocolEndpoint data.
PollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls.
Retries	0 to 10 retries Default: 3	Sets the number of retry polls to perform when the initial poll fails.
Timeout	10 to 10000 milliseconds Default: 700 milliseconds	Sets the amount of time to wait for the poll response before the first poll request times out. The timeout value doubles for each successive retry. For Timeout=700 msec (0.7 sec) and Retries=3: <ul style="list-style-type: none"> • 0.7 seconds for first retry • 1.4 seconds for second retry • 2.8 seconds for third retry

VRFs

The *VRFs* default polling group is used by the MPLS Manager to monitor the availability of VRF instances. It has a default setting named *VRF SNMP Setting* and an optional setting named *VRF External Setting*. The two settings are mutually exclusive because, as with most Domain Manager groups, only one setting can be applied to a group at any given time.

VRF SNMP Setting

The *VRF SNMP Setting* determines the polling intervals used by the MPLS Manager to monitor VRFs. MPLS Manager monitors the VRFs by probing the SNMP tables on the PE routers hosting the VRFs.

Table 33 lists the *VRF SNMP Setting* parameters.

Table 33: VRF SNMP Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: ENABLED	Enables or disables the MPLS Manager polling of VRFs and the subsequent analysis of polled VRF data.
PollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls.
Retries	0 to 10 retries Default: 3	Sets the number of retry polls to perform when the initial poll fails.
Timeout	10 to 10000 milliseconds Default: 700 milliseconds	Sets the amount of time to wait for the poll response before the first poll request times out. The timeout value doubles for each successive retry. For Timeout=700 msec (0.7 sec) and Retries=3: <ul style="list-style-type: none"> • 0.7 seconds for first retry • 1.4 seconds for second retry • 2.8 seconds for third retry

For the *VRF SNMP Setting*, MPLS Manager uses the values polled from the discovered PE routers when analyzing VRFs, including the values for MaxRoutes and MidRouteThreshold.

VRF External Setting

The *VRF External Setting* determines the MaxRoutes and MidRouteThreshold values used by the MPLS Manager when analyzing VRFs. (No SNMP polling is performed.) The analysis results are collected by an external process for test purposes.

Table 34 lists the *VRF External Setting* parameters.

Table 34: VRF External Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
MaxRoutes	Integer Default: 2000	Maximum number of routes that a VRF is configured to hold.
MidRoute Threshold	Integer Default: 100	Middle threshold for the number of routes that a VRF is configured to hold.

CE to CE Pings

The *CE to CE Pings* default remote ping group is used by the MPLS Manager to generate all possible CE-to-CE pings in a VPN, and to monitor and analyze the ping test results to determine CE-to-CE reachability. It has a single, valid setting named *CE to CE Ping Setting*.

This setting determines (1) the creation of remote ping entries in the SNMP Ping MIB tables on the source CEs and (2) the polling intervals used by the MPLS Manager to monitor the ping test results. The MPLS Manager monitors the ping test results by probing the Ping MIB tables on the source CEs.

By default, each CE in a VPN pings every other CE in the VPN.

Table 35 lists the *CE to CE Ping Setting* parameters.

Table 35: CE to CE Ping Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: DISABLED	Enables or disables the MPLS Manager triggering of CE-to-CE pings and the subsequent polling and analysis of the polled ping test results.
Classification	CE_TO_CE Default: CE_TO_CE	Sets the type of ping to execute: source CEs to destination CEs on a per VPN basis.
DestinationIP	Regular expression Default: *.*.*.* (By default, all loopback IP addresses of a destination CE are selected.)	Specifies the wildcard expression to use to select the loopback IP addresses of the destination CEs for CE-to-CE pings. An example entry is 163.*.*.*, which would limit the selected loopback IP addresses to just those beginning with 163.
DestinationName	Regular expression Default: * (By default, all CEs are selected as destination CEs.)	Specifies the wildcard expression to use to select the destination CEs for CE-to-CE pings. An example entry is C*, which would limit the selected destination CEs to just those having names beginning with C.
packetSize	0 to 65507 octets Default: 100 octets	Specifies the size of the IP packet to be transmitted in a ping request executed by the source CE.

Table 35: CE to CE Ping Setting Parameters and Their Values *(continued)*

PARAMETER	VALUES	DESCRIPTION
ping_TimeOut	1000 to 60000 milliseconds Default: 2000 milliseconds	Sets the amount of time for a source CE to wait for a ping response before the ping request times out.
pollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls to retrieve the ping test results from the source CE.
polling_TimeOut	2000 to 10000 milliseconds Default: 2000 milliseconds	Sets the amount of time to wait for a poll response before trying again to retrieve the ping test results from the source CE.
SourceName	Regular expression Default: * (By default, all CEs are selected as source CEs.)	Specifies the wildcard expression to use to select the source CEs for CE-to-CE pings. An example entry is <code>A*</code> , which would limit the selected source CEs to just those having names beginning with <code>A</code> .
SourceType	HUBS SPOKES BOTH Default: SPOKES	Specifies the source CEs for hub-and-spoke VPNs.

For each CE-to-CE remote ping configured by an administrator, the MPLS Manager writes the ping entry to the source CE and periodically refreshes the ping entry so that it does not expire until the administrator deliberately disables it, or until the MPLS Manager is shut down.

Each configured CE-to-CE remote ping is represented by the MPLS Manager as a RemotePing element that appears in the Topology Browser. For a description of the RemotePing element, see [Remote Ping Functionality](#) on page 97.

PE to CE Ping Setting

The *PE to CE Pings* default remote ping group is used by the MPLS Manager to generate all possible PE-to-CE pings in a VPN, and to monitor and analyze the ping test results to determine PE-to-CE reachability. It has a single, valid setting named *PE to CE Ping Setting*.

This setting determines (1) the creation of remote ping entries in the SNMP Ping MIB tables on the source PEs and (2) the polling intervals used by the MPLS Manager to monitor the ping test results. The MPLS Manager monitors the ping test results by probing the Ping MIB tables on the source PEs.

By default, each PE in a VPN pings every *remote* CE in the VPN. A remote CE is a CE attached to a peer PE.

Table 36 lists the *PE to CE Ping Setting* parameters.

Table 36: PE to CE Ping Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: DISABLED	Enables or disables the MPLS Manager triggering of PE-to-CE pings and the subsequent polling and analysis of the polled ping test results.
Classification	PE_TO_REMOTE_CE PE_TO_LOCAL_CE Default: PE_TO_REMOTE_CE	Sets the type of ping to execute: <ul style="list-style-type: none"> • Source PEs to destination Remote CEs (CEs attached to the peer PEs of a source PE) on a per VPN basis • Source PEs to destination Local CEs (CEs attached directly to a source PE) on a per VPN basis
DestinationIP	Regular expression Default: *.*.*.* (By default, all loopback IP addresses of a destination CE are selected.)	Specifies the wildcard expression to use to select the loopback IP addresses of the destination CEs for PE-to-CE pings. An example entry is 163.*.*.*, which would limit the selected loopback IP addresses to just those beginning with 163.

Table 36: PE to CE Ping Setting Parameters and Their Values *(continued)*

PARAMETER	VALUES	DESCRIPTION
DestinationName	Regular expression Default: * (By default, all CEs are selected as destination CEs.)	Specifies the wildcard expression to use to select the destination CEs for PE-to-CE pings. An example entry is C*, which would limit the selected destination CEs to just those having names beginning with C.
packetSize	0 to 65507 octets Default: 100 octets	Specifies the size of the IP packet to be transmitted in a ping request executed by the source PE.
ping_TimeOut	1000 to 60000 milliseconds Default: 2000 milliseconds	Sets the amount of time for a source PE to wait for a ping response before the ping request times out.
pollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls to retrieve the ping test results from the source PE.
polling_TimeOut	2000 to 10000 milliseconds Default: 2000 milliseconds	Sets the amount of time to wait for a poll response before trying again to retrieve the ping test results from the source PE.
SourceName	Regular expression Default: * (By default, all PEs are selected as source PEs.)	Specifies the wildcard expression to use to select the source PEs for PE-to-CE pings. An example entry is A*, which would limit the selected source PEs to just those having names beginning with A.
SourceType	HUBS SPOKES BOTH Default: SPOKES	Specifies the source PEs for hub-and-spoke VPNs.

For each PE-to-CE remote ping configured by an administrator, the MPLS Manager writes the ping entry to the source PE and periodically refreshes the ping entry so that it does not expire until the administrator deliberately disables it, or until the MPLS Manager is shut down.

Each configured PE-to-CE remote ping is represented by the MPLS Manager as a RemotePing element that appears in the Topology Browser. For a description of the RemotePing element, see [Remote Ping Functionality](#) on page 97.

PE to PE Ping Setting

The *PE to PE Pings* default remote ping group is used by the MPLS Manager to generate all possible PE-to-PE pings in a VPN, and to monitor and analyze the ping test results to determine PE-to-PE reachability. It has a single, valid setting named *PE to PE Ping Setting*.

This setting determines (1) the creation of remote ping entries in the SNMP Ping MIB tables on the source PEs and (2) the polling intervals used by the MPLS Manager to monitor the ping test results. The MPLS Manager monitors the ping test results by probing the Ping MIB tables on the source PEs.

By default, each PE in a VPN pings every other PE in the VPN.

Table 37 lists the *PE to PE Ping Setting* parameters.

Table 37: PE to PE Ping Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: DISABLED	Enables or disables the MPLS Manager triggering of PE-to-PE pings and the subsequent polling and analysis of the polled ping test results.
Classification	PE_TO_PE Default: PE_TO_PE	Sets the type of ping to execute: source PEs to destination PEs on a per VPN basis.
DestinationIP	Regular expression Default: *.*.*.* (By default, all loopback IP addresses of a destination PE are selected.)	Specifies the wildcard expression to use to select the loopback IP addresses of the destination PEs for PE-to-PE pings. An example entry is 163.*.*.*, which would limit the selected loopback IP addresses to just those beginning with 163.

Table 37: PE to PE Ping Setting Parameters and Their Values *(continued)*

PARAMETER	VALUES	DESCRIPTION
DestinationName	Regular expression Default: * (By default, all PEs are selected as destination PEs.)	Specifies the wildcard expression to use to select the destination PEs for PE-to-PE pings. An example entry is C*, which would limit the selected destination PEs to just those having names beginning with C.
packetSize	0 to 65507 octets Default: 100 octets	Specifies the size of the IP packet to be transmitted in a ping request executed by the source PE.
ping_TimeOut	1000 to 60000 milliseconds Default: 2000 milliseconds	Sets the amount of time for a source PE to wait for a ping response before the ping request times out.
pollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls to retrieve the ping test results from the source PE.
polling_TimeOut	2000 to 10000 milliseconds Default: 2000 milliseconds	Sets the amount of time to wait for a poll response before trying again to retrieve the ping test results from the source PE.
SourceName	Regular expression Default: * (By default, all PEs are selected as source PEs.)	Specifies the wildcard expression to use to select the source PEs for PE-to-PE pings. An example entry is A*, which would limit the selected source PEs to just those having names beginning with A.
SourceType	HUBS SPOKES BOTH Default: SPOKES	Specifies the source PEs for hub-and-spoke VPNs.

For each PE-to-PE remote ping configured by an administrator, the MPLS Manager writes the ping entry to the source PE and periodically refreshes the ping entry so that it does not expire until the administrator deliberately disables it, or until the MPLS Manager is shut down.

Each configured PE-to-PE remote ping is represented by the MPLS Manager as a RemotePing element that appears in the Topology Browser. For a description of the RemotePing element, see [Remote Ping Functionality](#) on page 97.

PE to Unmanaged CE Ping Settings

The *PE to Unmanaged CE Pings* default remote ping group is used by the MPLS Manager to generate all possible PE-to-remote-Unmanaged-CE pings in a VPN, and to monitor and analyze the ping test results to determine PE-to-Unmanaged-CE reachability. It has a single, valid setting named *PE to Unmanaged CE Ping Setting*.

This setting determines (1) the creation of remote ping entries in the SNMP Ping MIB tables on the source PEs and (2) the polling intervals used by the MPLS Manager to monitor the ping test results. The MPLS Manager monitors the ping test results by probing the Ping MIB tables on the source PEs.

By default, each PE in a VPN pings every *remote* Unmanaged CE in the VPN. A remote Unmanaged CE is an Unmanaged CE attached to a peer PE.

Note: The term *Unmanaged CE*, as used here, refers to a CE that is *not* in the repository of the MPLS Manager or the source Availability Manager(s).

Table 38 lists the *PE to Unmanaged CE Ping Setting* parameters.

Table 38: PE to Unmanaged CE Ping Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: DISABLED	Enables or disables the MPLS Manager triggering of PE-to-Unmanaged-CE pings and the subsequent polling and analysis of the polled ping test results.
Classification	PE_TO_UNMANAGED_CE Default: PE_TO_UNMANAGED_CE	Sets the type of ping to execute: source PEs to destination Unmanaged CEs on a per VPN basis.

Table 38: PE to Unmanaged CE Ping Setting Parameters and Their Values *(continued)*

PARAMETER	VALUES	DESCRIPTION
DestinationIP	Regular expression Default: *.*.*.* (By default, all VRF-associated interface IP addresses on a destination Unmanaged CE are selected; that is, all CE interface IP addresses associated with VRF interface IP addresses on a peer PE are selected.)	Specifies the wildcard expression to use to select the VRF-associated IP addresses on the destination Unmanaged CEs for PE-to-Unmanaged-CE pings. An example entry is 163.*.*.*, which would limit the selected VRF-associated IP addresses to just those beginning with 163.
DestinationName	Regular expression Default: * (By default, all Unmanaged CEs are selected as destination Unmanaged CEs.)	Specifies the wildcard expression to use to select the destination Unmanaged CEs for PE-to-Unmanaged-CE pings. An example entry is C*, which would limit the selected destination Unmanaged CEs to just those having names beginning with C.
packetSize	0 to 65507 octets Default: 100 octets	Specifies the size of the IP packet to be transmitted in a ping request executed by the source PE.
ping_TimeOut	1000 to 60000 milliseconds Default: 2000 milliseconds	Sets the amount of time for a source PE to wait for a ping response before the ping request times out.
pollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls to retrieve the ping test results from the source PE.
polling_TimeOut	2000 to 10000 milliseconds Default: 2000 milliseconds	Sets the amount of time to wait for a poll response before trying again to retrieve the ping test results from the source PE.
SourceName	Regular expression Default: * (By default, all PEs are selected as source PEs.)	Specifies the wildcard expression to use to select the source PEs for PE-to-Unmanaged-CE pings. An example entry is A*, which would limit the selected source PEs to just those having names beginning with A.
SourceType	HUBS SPOKES BOTH Default: SPOKES	Specifies the source PEs for hub-and-spoke VPNs.

For each PE-to-Unmanaged-CE remote ping configured by an administrator, the MPLS Manager writes the ping entry to the source PE and periodically refreshes the ping entry so that it does not expire until the administrator deliberately disables it, or until the MPLS Manager is shut down.

Each configured PE-to-Unmanaged-CE remote ping is represented by the MPLS Manager as a RemotePing element that appears in the Topology Browser. For a description of the RemotePing element, see [Remote Ping Functionality](#) on page 97.

PE to VRF Ping Setting

The *PE to VRF Pings* default remote ping group is used by the MPLS Manager to generate all possible PE-to-remote-VRF pings in a VPN, and to monitor and analyze the ping test results to determine PE-to-VRF reachability. It has a single, valid setting named *PE to VRF Ping Setting*.

This setting determines (1) the creation of remote ping entries in the SNMP Ping MIB tables on the source PEs and (2) the polling intervals used by the MPLS Manager to monitor the ping test results. The MPLS Manager monitors the ping test results by probing the Ping MIB tables on the source PEs.

By default, each PE in a VPN pings every *remote* VRF in the VPN. A remote VRF is a VRF hosted by a peer PE.

Table 39 lists the *PE to VRF Ping Setting* parameters.

Table 39: PE to VRF Ping Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: DISABLED	Enables or disables the MPLS Manager triggering of PE-to-VRF pings and the subsequent polling and analysis of the polled ping test results.
Classification	PE_TO_VRF Default: PE_TO_VRF	Sets the type of ping to execute: source PEs to destination VRFs on a per VPN basis.

Table 39: PE to VRF Ping Setting Parameters and Their Values *(continued)*

PARAMETER	VALUES	DESCRIPTION
DestinationIP	Regular expression Default: *.*.*.* (By default, all interface IP addresses associated with the destination VRF are selected.)	Specifies the wildcard expression to use to select the interface IP addresses associated with the destination VRFs for PE-to-VRF pings. An example entry is 163.*.*.*, which would limit the selected interface IP addresses to just those beginning with 163.
DestinationName	Regular expression Default: * (By default, all destination VRFs are selected.)	Specifies the wildcard expression to use to select the destination VRFs for PE-to-VRF pings. An example entry is C*, which would limit the selected destination VRFs to just those having names beginning with C.
packetSize	0 to 65507 octets Default: 100 octets	Specifies the size of the IP packet to be transmitted in a ping request executed by the source PE.
ping_TimeOut	1000 to 60000 milliseconds Default: 2000 milliseconds	Sets the amount of time for a source PE to wait for a ping response before the ping request times out.
pollingInterval	30 to 3600 seconds Default: 240 seconds	Sets the time between successive polls to retrieve the ping test results from the source PE.
polling_TimeOut	2000 to 10000 milliseconds Default: 2000 milliseconds	Sets the amount of time to wait for a poll response before trying again to retrieve the ping test results from the source PE.
SourceName	Regular expression Default: * (By default, all PEs are selected as source PEs.)	Specifies the wildcard expression to use to select the source PEs for PE-to-VRF pings. An example entry is A*, which would limit the selected source PEs to just those having names beginning with A.
SourceType	HUBS SPOKES BOTH Default: SPOKES	Specifies the source PEs for hub-and-spoke VPNs.

For each PE-to-VRF remote ping configured by an administrator, the MPLS Manager writes the ping entry to the source PE and periodically refreshes the ping entry so that it does not expire until the administrator deliberately disables it, or until the MPLS Manager is shut down.

Each configured PE-to-VRF remote ping is represented by the MPLS Manager as a RemotePing element that appears in the Topology Browser. For a description of the RemotePing element, see [Remote Ping Functionality](#) on page 97.

System Write Community Strings

The *System Write Community Strings* default remote ping group is used by periodic remote pings (deployed by the remote ping groups *CE to CE Pings* through *PE to VRF Pings*) and by on-demand remote pings to write remote ping entries to the Ping-MIB-enabled systems in the managed MPLS environment. It has a single setting named *Router Community Strings*.

Table 40 lists the *Router Community Strings* setting parameters.

Table 40: Router Community Strings Setting Parameters

PARAMETER	VALUES	DESCRIPTION
WriteCommunity	String of unspecified length Default: "private"	The write community string for the router(s) defined for this group.

By default, all discovered routers in the managed MPLS environment become members of the *System Write Community Strings* group. Thus, by default, for both periodic and on-demand remote pings, the MPLS Manager uses the write community string (also known as the write community name) defined in the *System Write Community Strings* group for all Ping-MIB-enabled routers in the managed MPLS environment.

If individual routers or subgroups of routers in the managed MPLS environment use different write community strings, an administration must create a new *System Write Community Strings* group for each write community string. The basic steps are:

- 1 Create a *System Write Community Strings* group for each router or subgroup of routers having a different write community string; for instructions on creating new groups, see [Creating New Groups](#) on page 94.

- 2 Assign a different write community string to each of the *System Write Community Strings* groups; for instructions on modifying a setting's parameters, see [Modifying the Parameters of a Setting](#) on page 93.
- 3 Define matching criteria for each of the *System Write Community Strings* groups to limit each group's membership to the router or subgroup of routers to which the write community string applies; for instructions on defining matching criteria, see [Editing Matching Criteria](#) on page 92.
- 4 Change the priority of the *System Write Community Strings* groups if need be; for instructions on prioritizing groups, see [Modifying the Priority of Groups](#) on page 91.

To write a ping entry to a router, the MPLS Manager searches through the prioritized list of *System Write Community Strings* groups for a router match, finds a match, includes the associated write community string in the remote ping request, and sends the request to the router. If no match is found, the MPLS Manager includes the global write community string defined in the *REMOTEPING.conf* file in the remote ping request. For a description of the *REMOTEPING.conf* file, see the *EMC Smarts MPLS Manager Configuration Guide*.

Default Threshold Groups and Settings

Currently, there are no default threshold groups and settings for MPLS Manager. The MPLS Manager compares the polled data against internal, fixed thresholds to determine fault conditions.

Opening the Polling and Thresholds Console

The Polling and Thresholds Console is used to display groups and modify their properties. To access the Polling and Threshold Console, you must first open the Domain Manager Administration Console.

Attaching to a Domain Manager, such as the MPLS Manager, with the Domain Manager Administration Console requires an InCharge user account with the following privileges and permissions:

- All privileges, specified in the *serverConnect.conf* file (or its equivalent) read by the Domain Manager.

- Permission to use the console operation named *Configure Domain Manager Admin Console*. Through the Global Manager Administration Console, this permission is specified in the *Console Operations* section of the user profile.

For information about configuring access privileges, see the *InCharge System Administration Guide*. For information about configuring permissions to perform specific console operations, see the *EMC Smarts Service Assurance Manager Configuration Guide*.

To open the Polling and Thresholds Console, follow these steps:

- 1** Attach the Global Console to the Domain Manager. The Topology Browser Console opens.
- 2** In the Topology Browser Console, select *Configure > Domain Manager Administration Console*. The Domain Manager Administration Console opens.
- 3** In the Domain Manager Administration Console, select *Edit > Polling and Thresholds*. The Polling and Thresholds Console opens.

Layout of the Polling and Thresholds Console

The Polling and Thresholds Console is divided into two panels.

- The left panel displays the icon for the analysis domain in the upper-left corner and provides two tabs, Polling and Thresholds, at the bottom. When the Polling tab is selected, the console displays polling groups (and, for the MPLS Manager, also displays remote ping groups). Likewise, when the Thresholds tab is selected, the console displays threshold groups. (Threshold groups are not available to the MPLS Manager.)

For each group, there are settings that provide adjustable parameters and a membership list of managed elements to which the settings are applied.

- The right panel remains blank until a group, setting, or member is selected in the left panel. When an item is selected in the left panel, the right panel displays additional information regarding that item.

Figure 27 provides an example of a Polling and Thresholds Console. attached to an MPLS Manager named `INCHARGE-MPLS`.

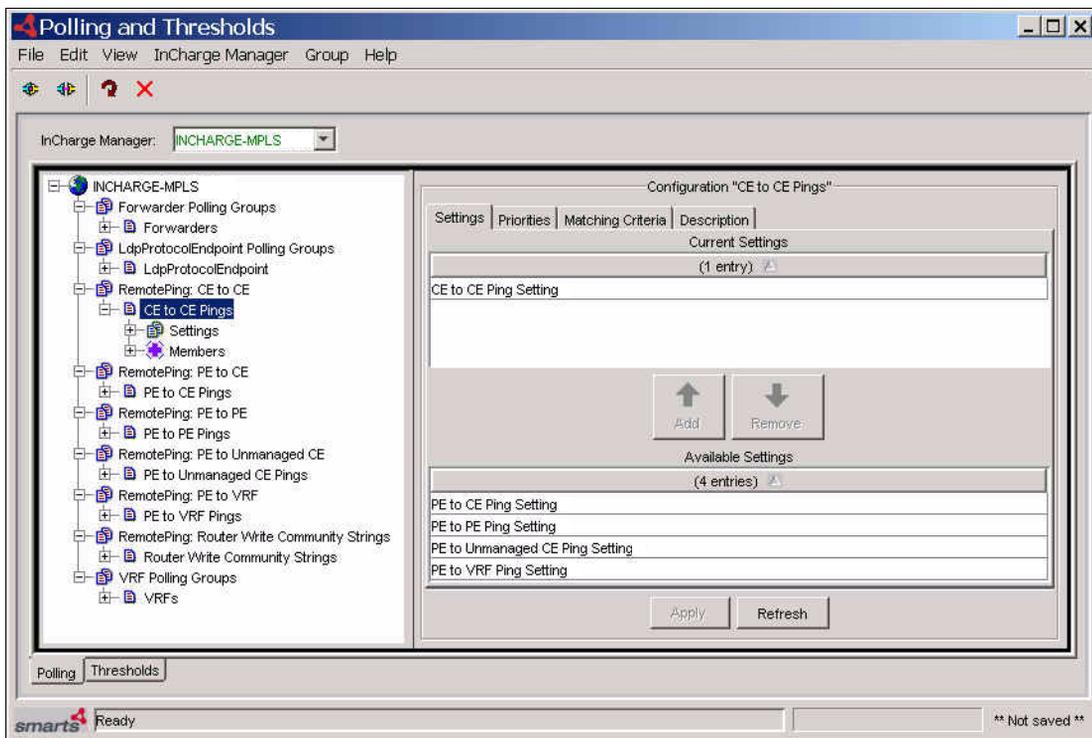


Figure 27: Polling and Thresholds Console—Example

Polling and Thresholds Console Toolbar Buttons

The toolbar of the Polling and Thresholds Console provides quick access to the commands described in Table 41.

Table 41: Polling and Thresholds Console Toolbar Buttons

BUTTON	DESCRIPTION
	Attach to a Domain Manager
	Detach from a Domain Manager
	Reconfigure polling and thresholds groups
	Delete selected item

Working With Groups and Settings

A group is composed of settings and members. A setting is composed of one or more related parameters, and a member is an element of the managed topology that belongs to a group.

Using the Polling and Thresholds Console, you can perform the following configuration tasks:

- Modify the properties of existing groups.
 - Determine what setting is applied to a group.
 - Modify the parameters of a setting.
- Create new groups.

How Managed Elements Are Assigned to Groups

When a Domain Manager performs discovery, it automatically assigns each managed element to a group based on:

- Matching criteria defined for the group
- Priority of the group, which determines membership when a device meets the matching criteria for more than one group

A managed element can be a member of one and only one group.

Modifying the Properties of a Group

A group is composed of settings and members. A setting includes one or more polling parameters. The matching criteria specified for the group and the group's priority determine which managed elements become members of the group.

When a group name is selected in the left panel of the Polling and Thresholds Console, four tabs are displayed:

- Settings
- Priorities
- Matching Criteria
- Description

Modifying the properties under each of these tabs changes the configuration of the group. When you finish editing the properties of a group, click the **Apply** button to save the changes and then select **Reconfigure** from the *Group* menu to make the configuration changes take effect.

Adding or Removing Group Settings

A group's applied setting or settings determine what parameters are applied to the managed elements that are members of the group. For all Domain Managers except the Performance Manager, only one setting may be applied to a group at any given time. For the Performance Manager, one and only one connectivity type setting plus zero or more performance type settings may be applied simultaneously to a group.

The Settings tab is divided into two sections: Current Settings and Available Settings. The Current Settings section lists the setting that is applied to the group. The Available Settings section lists additional available settings.

Note: For the MPLS Manager, do not change the current setting for any of the remote ping groups.

Adding or Removing a Setting

- 1 Select a setting from the Current Settings list or from the Available Settings list.
- 2 Click **Add** to move an available setting to the Current Settings list or click **Remove** to move a current setting to the Available Settings list.
- 3 Click **Apply**.
- 4 Select **Reconfigure** from the *Group* menu.

Modifying the Priority of Groups

Matching criteria and priority determine which managed elements are members of what group. When an element matches the criteria for two or more groups, the managed element becomes a member of the group with the highest priority. The Priorities tab lists groups in the order of their priority, from highest to lowest.

Changing the Priority of a Group

- 1 Select the group for which you want to change the priority.
- 2 Click the up or down arrow to change its position relative to the other groups.
- 3 Click **Apply**.
- 4 Select **Reconfigure** from the *Group* menu.

Editing Matching Criteria

Matching criteria, which appear at the top of the Matching Criteria tab, are defined using the attributes of the matching-criteria element type associated with the group; for example, VRF attributes for MPLS Manager's VRFs polling group. Each matching criterion has three fields: *Name*, *Description*, and *Value*.

- Name identifies the attribute that is used as a matching criterion.
- Description is the description of the attribute taken from the ICIM model.
- Value is any combination of text, integers, and wildcards (see [Wildcard Patterns](#) on page 129) that is matched against the value of the attribute in the managed element. The *value* field for a matching criterion is *not* case-sensitive.

If the value of a managed element's attribute matches a matching criterion, the managed element is eligible to become a member of the group. For example, if a matching criterion uses the attribute *SystemName* with a value of **30**, all members of the group must contain the string *30* in their *SystemName* attribute. When more than one matching criterion is specified, a managed element must match all criteria to become a member of the group.

Adding or Removing Matching Criteria

- 1 Select a matching criterion.
- 2 Click **Enable** to make the criterion active, moving it to the top of the Matching Criteria tab.

(Use **Disable** to deactivate the criterion, moving it to the bottom of the Matching Criteria tab.)
- 3 If you are adding a matching criterion, type a matching pattern in the *value* field.

- 4 Click **Apply**.
- 5 Select **Reconfigure** from the *Group* menu.

Changing the Value of a Matching Criterion

- 1 Select the string in the `value` field or double-click the `value` field to highlight the current value.
- 2 Type the text, integers, or wildcard to match against the attribute.
- 3 Click **Apply**.
- 4 Select **Reconfigure** from the *Group* menu.

A Domain Manager processes matching criteria in the following manner. First, managed elements are compared against the matching criteria of the group with the highest priority. If an element matches all the criteria, it is added as a member of the group. If an element does not match all the criteria, it is compared against the matching criteria of the group with the second highest priority, and so on.

When no matching criteria are active (or appear in the top of the Matching Criteria dialog box), the group matches all managed elements of the matching-criteria element type associated with the group.

Modifying the Parameters of a Setting

The parameters of a setting are changed in one of two ways: by (1) choosing a value from a drop-down menu or (2) entering a value in a `value` field or adjusting a slider bar representing a range of values.

Changing the Parameters of a Setting

- 1 Select the setting in the left panel of the Polling and Thresholds Console. The parameters of a setting are listed in the right panel of the console.
- 2 Change the value of a parameter using one of the following methods:
For a drop-down menu, click the menu and select a value.
For a slider bar presentation,
 - Type a value into the `value` field and press **Enter** or
 - Select the slider bar and drag its handle with the mouse to change the value or select the slider bar and use the arrow keys to move its handle to change the value.
- 3 Click **Apply** to save the changes.
- 4 Select **Reconfigure** from the *Group* menu.

Restoring the Default Values of a Setting

The **Restore Defaults** button, which is visible when a setting is selected in the left panel of the Polling and Thresholds Console, restores the default values of all the parameters for the selected setting.

- 1 Select the setting.
- 2 Click **Restore Defaults**.
- 3 Select **Reconfigure** from the *Group* menu.

Creating New Groups

Creating a new group enables you to customize the settings for a group of managed elements. You can use two methods to create a new group:

- Copy an existing group. The new group contains the same settings and thresholds as the original group. Matching criteria are not copied.
- Create an empty group. The new group does not contain any settings or members. You must add settings and matching criteria, and set the priority of the new group.

After you create a new group, use procedures previously described to adjust the settings of the new group. For information regarding settings, see [Modifying the Priority of Groups](#) on page 91, and for information regarding groups, see [Modifying the Properties of a Group](#) on page 90.

Copying an Existing Group

- 1 Right-click the Polling or Threshold group that you want to copy.
- 2 Select **Copy** from the pop-up menu to display the Copy Group dialog.
- 3 In the dialog, type a name and an optional description for the new group and click **OK**. The new group contains the same settings and thresholds as the group you copied.
- 4 Edit the settings, matching criteria, and priority of the new group. Change the value of any parameters as necessary.
- 5 Select **Reconfigure** from the *Group* menu.

Creating an Empty Group

- 1** In the left panel of the Polling and Threshold Console, right-click the group type for which you want to create a new group. (When a Domain Manager provides more than one default group, you can create more than one type of group.)
- 2** Select **New Group** from the pop-up menu to display the New Group dialog.
- 3** In the dialog, type a name and an optional description for the new group and click **OK**.
- 4** Add settings and matching criteria, and set the priority of the new group. Change the values of any parameters as necessary.
- 5** Select **Reconfigure** from the *Group* menu.

6

Remote Ping Functionality

Remote ping allows MPLS network elements (PE and CE devices) to ping one another (and associated VRFs) to get an indication of the customer experience, and to determine the reachability of network elements from various network devices. *Remote ping is only available for L3VPNs.*

The MPLS Manager uses Internet Control Message Protocol (ICMP) Echo (ping requests), as documented in RFC 792. These remote ping requests may be configured to be sent from source to destination devices within the managed MPLS network and to certain CE devices outside the managed network.

Types of Remote Ping Requests

There are two types of remote ping requests:

- Periodic—Automated remote ping requests, configured (generally, by an administrator) through the Polling and Thresholds Console to run at configured intervals for an indefinite period of time.
- On-demand—Manual remote ping requests, launched from a source device to a destination device. On-demand requests are available from the right-click server tools menu for elements such as routers and VRFs displayed in the Notification Log, Topology Browser, and MPLS maps.

For periodic remote ping requests, the MPLS Manager analyzes the ping responses and notifies the Global Manager when the number of failed pings within configured polling interval meet or exceed the value configured in the *RemotePing.conf* file. For on-demand ping requests, the MPLS Manager returns the ping responses to the initiator of the ping requests but does not analyze the responses.

How Remote Pings are Generated

Table 42 describes how the different kinds of remote pings are generated.

Table 42: How Remote Pings are Generated

TYPE	DESCRIPTION
PE to PE	A PE (source PE) in a VPN pings another PE (destination PE) in the VPN.
CE to CE	A CE (source CE) in a VPN pings another CE (destination CE) in the VPN.
PE to remote CE	A PE (source PE) in a VPN pings a remote CE (destination CE) in the VPN. A remote CE is a CE attached to a peer PE of the source PE.
PE to local CE	A PE (source PE) in a VPN pings a local CE (destination CE) in the VPN. A local CE is a CE attached directly to the source PE.
PE to unmanaged CE	A PE (source PE) in a VPN pings a remote Unmanaged CE (destination CE) in the VPN. A remote Unmanaged CE is an Unmanaged CE attached to a peer PE of the source PE.
PE to VRF	A PE (source PE) in a VPN pings a remote VRF (destination VRF) in the VPN. A remote VRF is a VRF hosted by a peer PE of the source PE.

Examples of Remote Pings

Remote Ping Example 1: PE to PE

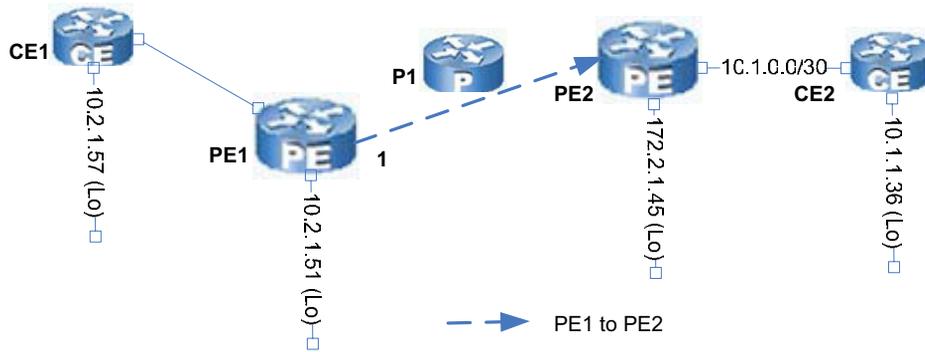


Figure 28: Example 1—PE to PE

Remote Ping Example 2: CE to CE

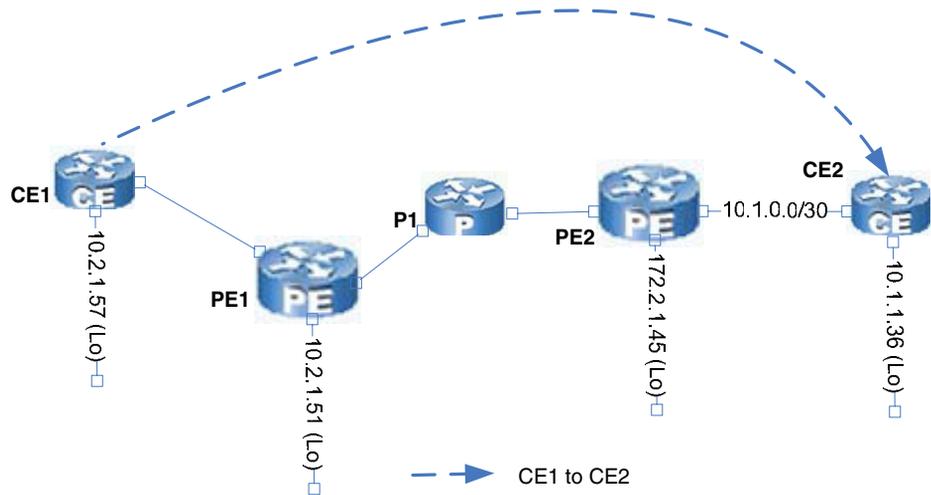


Figure 29: Example 2—CE to CE

Remote Ping Example 3: PE to Remote CE

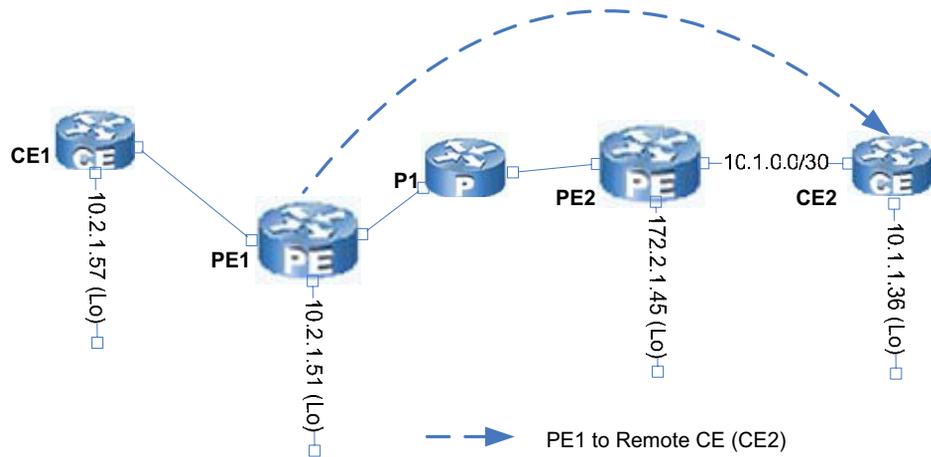


Figure 30: Example 3—PE to Remote CE

Remote Ping Example 4: PE to Local CE

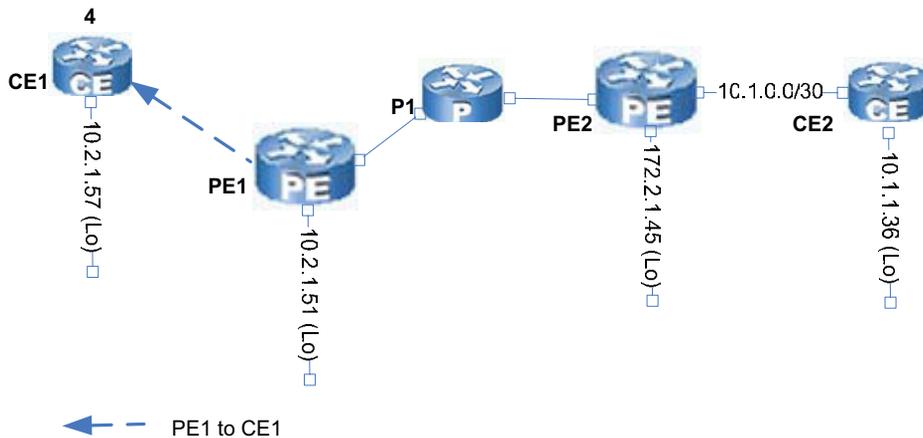


Figure 31: Example 4—PE to Local CE

Remote Ping Example 5: PE to Unmanaged CE

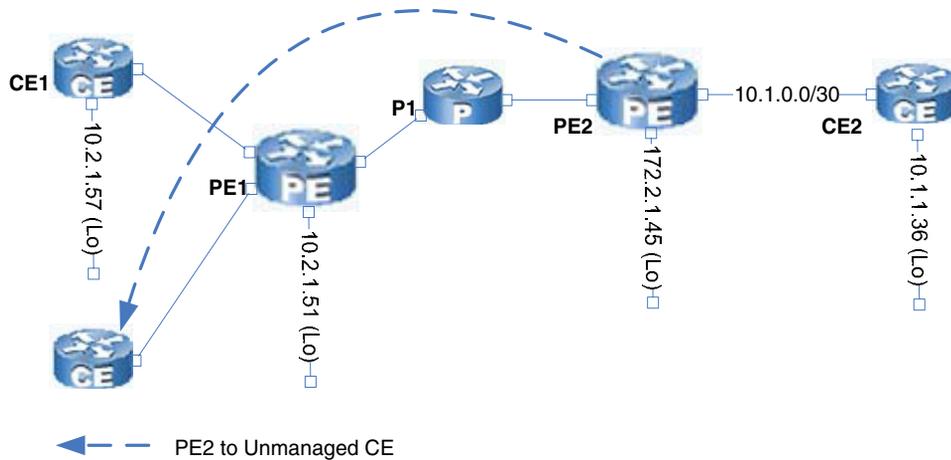


Figure 32: Example 5—PE to Unmanaged CE

Remote Ping Example 6: PE to VRF

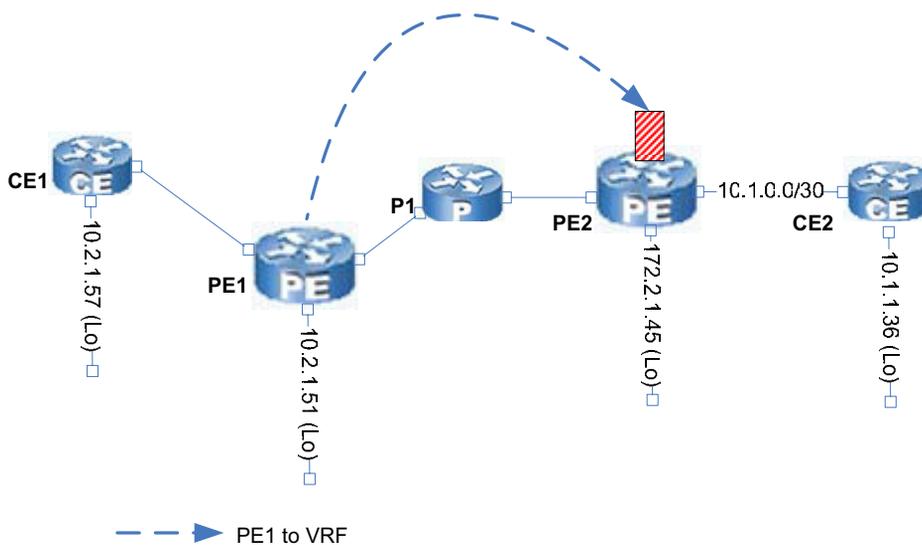


Figure 33: Example 6—PE to VRF

More about VRF Ping

When a VRF is the destination, the MPLS Manager pings all of the IP addresses of the underlying interfaces that are associated with the destination VRF. To determine which VRF to use, the MPLS Manager traces the topology to determine which VRF is used for a particular route. For example: Figure 34 shows two PE to VRF paths. In this case, when the MPLS Manager initiates Remote Ping 1 from CE1 to CE2, it must first determine which VRF handles the path to CE2. If multiple VRFs on a PE exchange routes with another VRF, the MPLS Manager pings all of those VRFs.

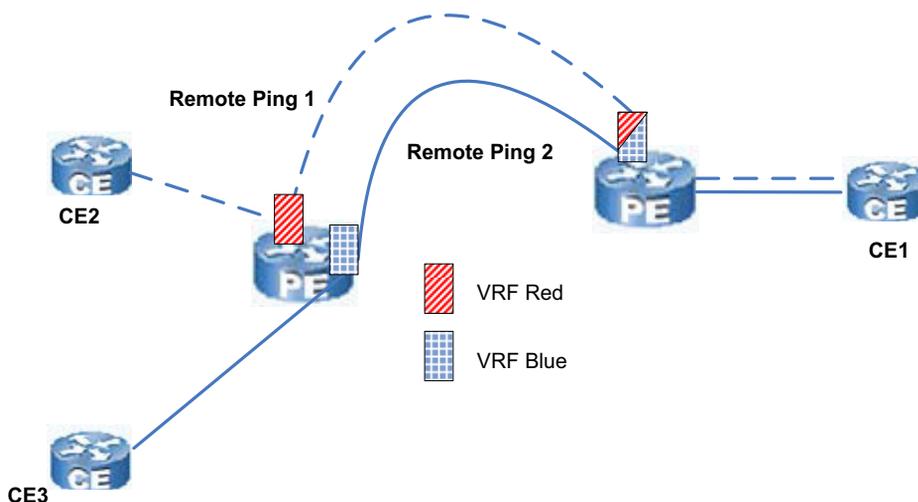


Figure 34: VRF-Aware Remote Ping—Example

Remote Ping Elements

A periodic remote ping instance is represented as a logical instance of the RemotePing ICIM class, that is, represented as a RemotePing element (object).

The MPLS Manager builds a data model of the periodic RemotePing elements as they are configured at the Polling and Thresholds Console. Each RemotePing object is created with certain attributes and relationships, and can experience certain events. For information about configuring periodic remote pings, see [Customizing Groups and Settings](#) on page 69.

Whenever a user configures a periodic remote ping instance, the MPLS Manager creates a RemotePing instance in the MPLS/VPN topology. RemotePing elements are displayed in a Topology Browser attached to the MPLS Manager. When analysis indicates a problem with a periodic remote ping, that RemotePing element is notified in the Notification Log.

Attributes for RemotePing Elements

Table 43 lists the attributes for RemotePing.

Table 43: Attributes for RemotePing Elements

ATTRIBUTE	DESCRIPTION
Name	The name of the remote ping element. The name is of the form: RP- <i><source_ip></i> - <i><destination_ip></i> - <i><packet_size></i> - <i><requestdelay></i> - <i><# of request></i> - <i><timeout></i> - <i><source_vrfname></i>
TimeStarted	The time that the remote ping was first started.
TimeRefreshed	The last time that the remote ping was refreshed.
Key	The index of the entry in the ping MIB on the source device that corresponds to this remote ping.
PacketsSent	The number of ICMP Echo requests sent during the last polling interval.
PacketsReceived	The number of ICMP Echo replies received during the last polling interval.
Description	A description of the RemotePing object. The description is of the form: <i><ping_type></i> from <i><source></i> to <i><destination></i>

Relationships for RemotePing Elements

Table 44 lists the relationships for RemotePing.

Table 44: Relationships for RemotePing Elements

RELATIONSHIP	Description
ConnectedSystems	The routing devices used as the source and destination of this remote ping.
Underlying	The VPN with which this remote ping is associated.

Remote Ping Impact Analysis

The MPLS Manager performs impact analysis on the RemotePing and L3VPN VPN managed elements and notifies the results in the Notification Log.

Table 45 lists the impact events notified by the MPLS Manager for the RemotePing and L3VPN elements, including the condition for each event.

Table 45: RemotePing Impact Events Notified by MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
RemotePing	Impaired	Caused by one of the following: <ul style="list-style-type: none"> • Underlying RemotePing Impaired event • Performance Manager event • Flapping event on underlying interfaces notified by Availability Manager
	Down	Caused by one of the following: <ul style="list-style-type: none"> • An underlying RemotePing Down event • Down event on underlying interfaces notified by Availability Manager • Down event on underlying destination routing device (hosts the destination IP address)
VPN (L3VPN only)	ImpactedByRemotePingFailure	Caused by one of the following: <ul style="list-style-type: none"> • RemotePing Impaired • RemotePing Down

Viewing Periodic RemotePing Information

You can view periodic RemotePing information in the following Global Console locations:

- Notification Log—(Global Console attached to the Global Manager) Displays notifications for unsuccessful periodic remote pings, including Impaired and Down notifications; see [Notification Log](#) on page 105.
- Notification Properties—(Global Console attached to the Global Manager) Displays detailed information about each periodic remote ping listed in the Notification Log; see [Notification Properties Window](#) on page 106.
- Topology Browser—(Global Console attached to the MPLS Manager) Displays configured RemotePing elements for the MPLS domain, including result data in the Attributes tab; see [Topology Browser](#) on page 107.

For detailed description of Global Console functionality, see *EMC Smarts Service Assurance Manager Operator's Guide*.

Notification Log

You can view notifications for unsuccessful remote pings in the Notification Log. For each notified RemotePing element, the log shows the attributes, including the impact analysis results. Figure 35 shows a sample Notification Log with a RemotePing Impaired notification.

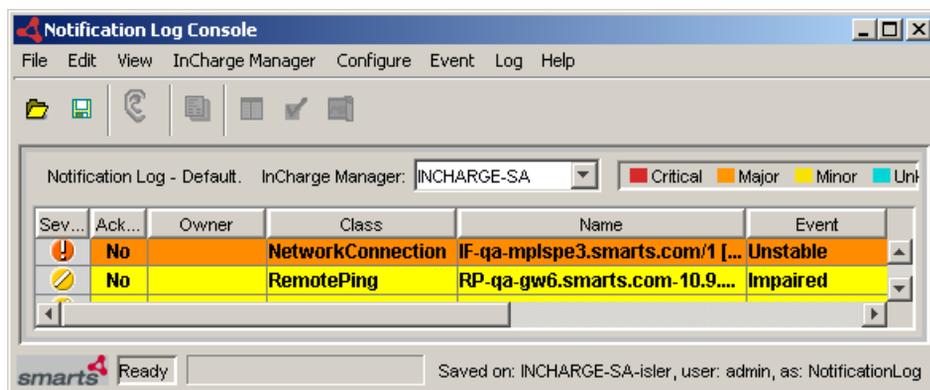


Figure 35: Notification Log Showing a RemotePing Notification

Notification Properties Window

You can see detailed information about a RemotePing by looking at the Notification Properties window. Figure 36 shows the Details tab for a RemotePing Down notification. The Details tab provides details about packet failures, packets sent, packets received, the source and destination elements, and so on.

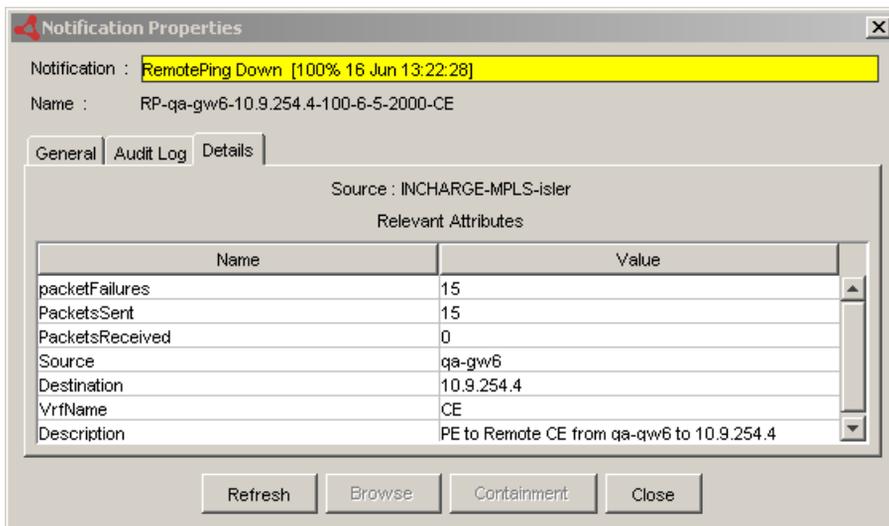


Figure 36: Notification Properties Window Showing RemotePing Details

Topology Browser

A Topology Browser attached to the MPLS Manager includes any configured periodic RemotePing objects in the topology tree for the MPLS Domain. The Attributes tab displays detailed result information. Figure 37 shows an example Topology Browser.

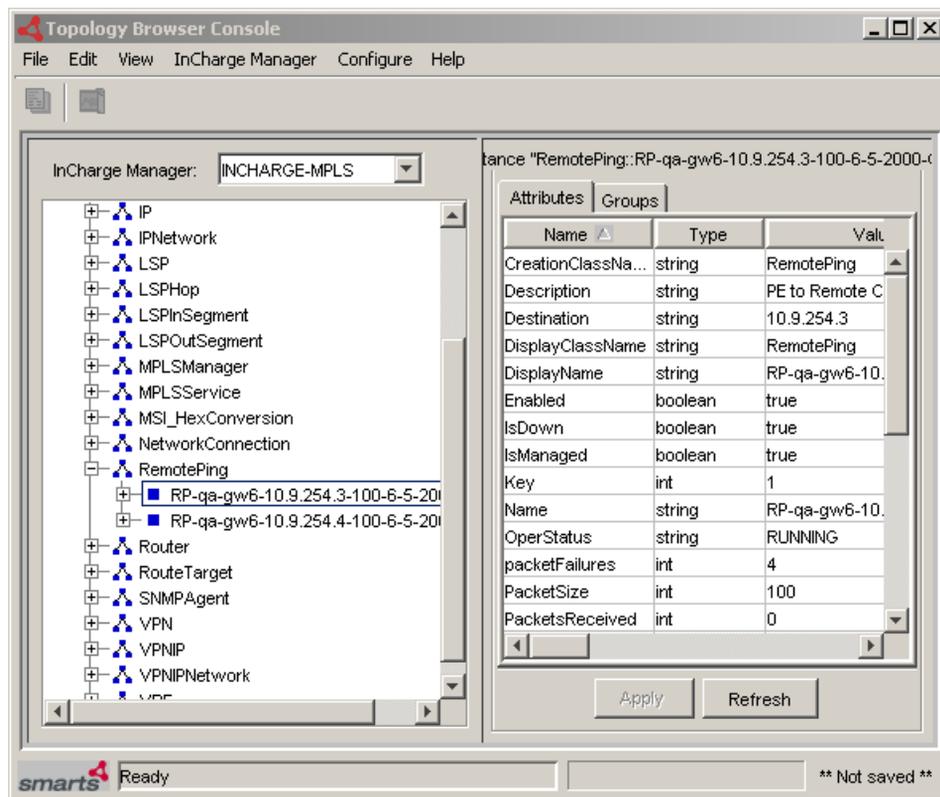


Figure 37: Topology Browser Showing RemotePing Elements

Issuing an On-Demand Remote Ping

With the Global Console attached to the Global Manager, you issue an on-demand remote ping for a particular routing device using one of the server tools listed in Table 46, available through the right-click menu for that routing device in the Topology Browser, Notification Log, or Map display.

Table 46: Remote Ping Server Tools

MENU OPTION	DESCRIPTION
Set Ping Source	Use this server tool to set a routing device as the source for on-demand remote pings. Once a source tool is set, you can use the device as the source for the Remote Ping and VRF Ping server tools.
Who's My Ping Source	Use this server tool to determine the current remote ping source.
Remote Ping	Use this server tool to issue a remote ping to PE and CE devices from the source device you have set using the Set Ping Source server tool to a device or devices. You can select the option of including all IP addresses (based on filter criteria) hosted by the destination device or just its SNMPAgent address.
VRF Ping	Use this server tool to issue a remote ping to one or more VRFs in the managed MPLS network.
Repeat Remote Ping	Use this server tool to rerun a periodic RemotePing that is notified in the Notification Log as Impaired or Down. This tool immediately sends five packets from the source device to the destination device with no delay between the packets.

Using the Set Ping Source Server Tool

To set a device as your ping source, follow these steps:

- 1** Connect the Global Console to the Global Manager.
- 2** On a Topology Browser, Notification Log, or Map display, select the PE or CE element to set as the source; for example, a PE named *qa-gw6.smarts.com*.
- 3** Right click the element, and then select *Server Tools*.
- 4** At the Server Tools menu, select *Set Ping Source*.

The results are displayed in the RemotePing - Set Ping Source text box, as shown in Figure 38.

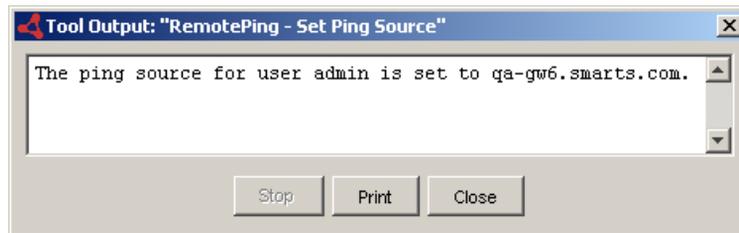


Figure 38: Set Ping Source Text Box

Using the Who's My Ping Source Server Tool

To determine which device is currently set as your ping source, follow these steps:

- 1 On a Topology Browser, Notification Log, or Map display, select the PE or CE element whose ping source you want to determine; for example, a PE named *qa-mplspe3.smarts.com*.
- 2 Right click the element, and then select *Server Tools*.
- 3 At the Server Tools menu, select *Who's My Ping Source*.

The results are displayed in the RemotePing - Who's My Ping Source text box, as shown in Figure 39. In this example, the ping source for *qa-mplspe3.smarts.com* is *qa-gw6.smarts.com*. In the text, user *admin* indicates that the user named "admin" issued the ping using the server tool.



Figure 39: Who's My Ping Source Text Box

Using the Remote Ping Server Tool

To launch an on-demand remote ping from your previously set ping source to a destination PE or CE routing device, follow these steps:

- 1 On a Topology Browser, Notification Log, or Map display, select the PE or CE element that is to be the destination; for example, a PE named *qa-gw3.smarts.com*.
- 2 Right click the element, and then select *Server Tools*.
- 3 At the Server Tools menu, select *Execute Ping Source*.

The results are displayed in the RemotePing - Execute Ping Source text box, as shown in Figure 40.

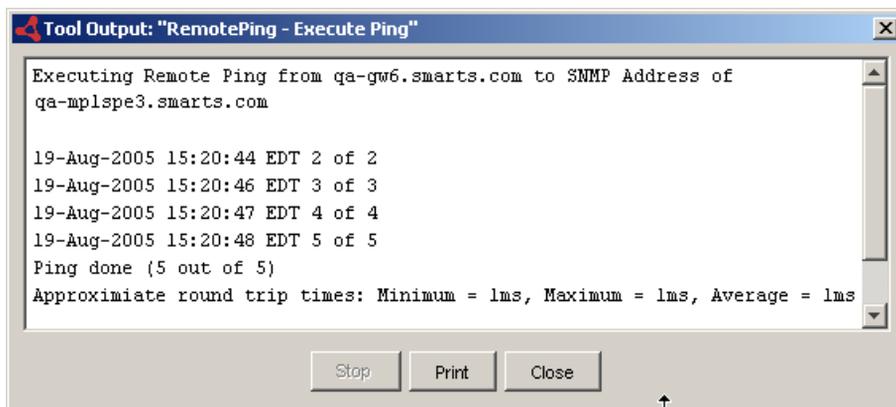


Figure 40: Execute Ping Text Box

Using the VRF Ping Server Tool

To launch an on-demand remote ping from your previously set ping source to a destination VRF, follow these steps:

- 1 On a Topology Browser, Notification Log, or Map display, select the VRF that is to be the destination.
- 2 Right click the VRF node, and then select *Server Tools*.
- 3 At the Server Tools menu, select *Execute VRF Ping*.

The results are displayed in the RemotePing - Execute VRF Ping text box, as shown in Figure 41.

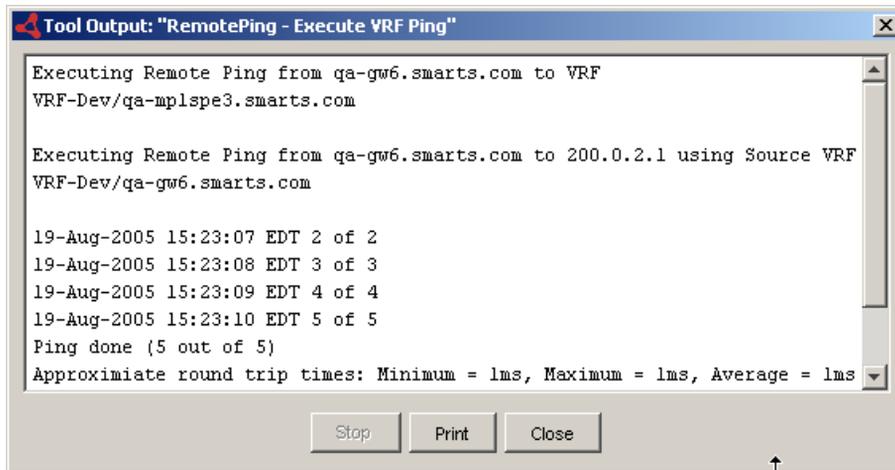


Figure 41: Execute VRF Ping Text Box

Using the Repeat Remote Ping Server Tool

To repeat a periodic remote ping for which you saw a notification of Impaired or Down, follow these steps:

- 1 At a Notification Log, select the notification for the periodic RemotePing you want to repeat.
- 2 Right click the notification, and then select *Server Tools*.
- 3 At the Server Tools menu, select *RemotePing-Repeat Remote Ping*.

The results are displayed in the RemotePing Execute VRF Ping text box, as shown in Figure 42.

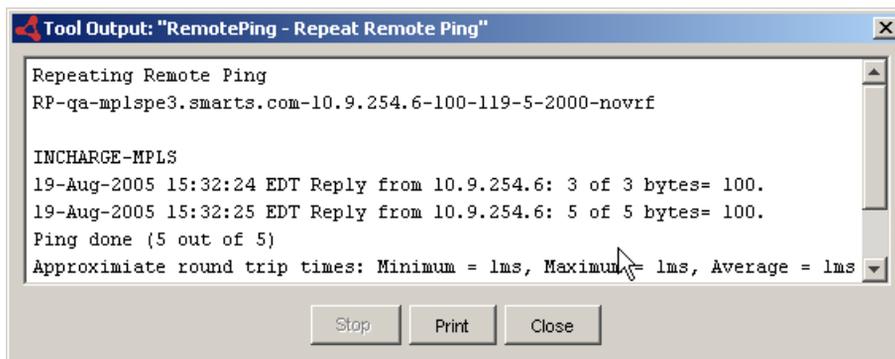


Figure 42: Repeat Remote Ping Text Box

Log Files

Log files for remote ping activity are located in **BASEDIR**/*smarts/logs/INCHARGE-MPLS.log* file in the MPLS Manager installation directory.

MPLS Terminology

The terms and concepts presented in this appendix should prove helpful in understanding the MPLS and VPN elements discovered and monitored by the MPLS Manager.

Begin by examining the following diagram:

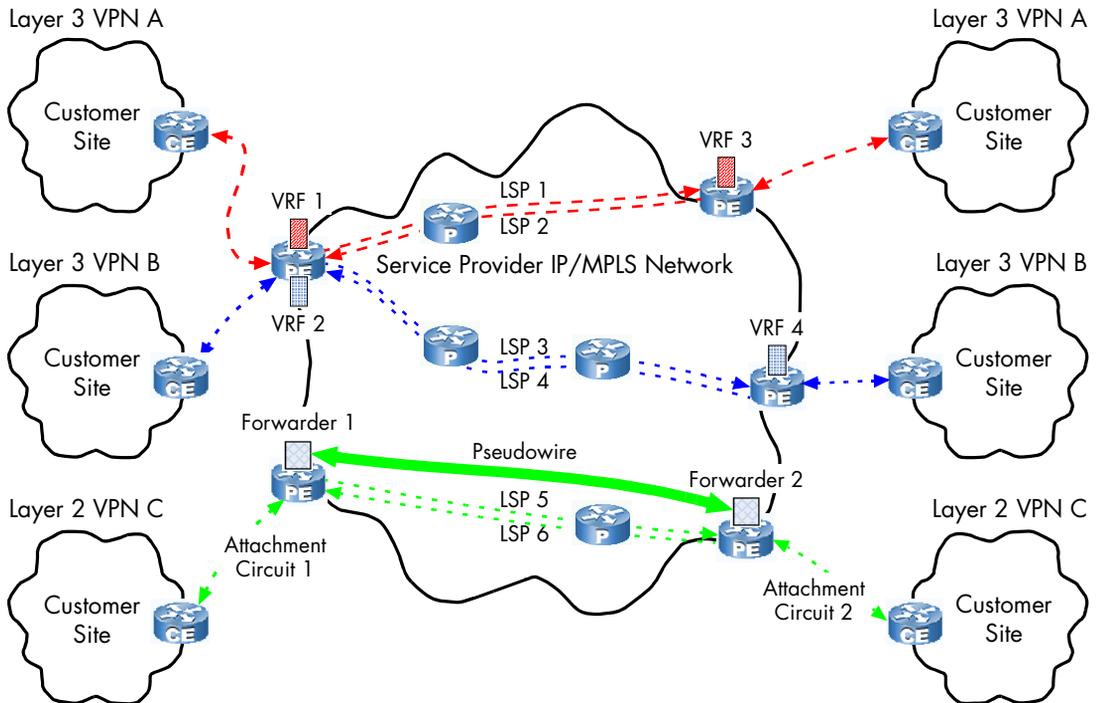


Figure 43: MPLS-Capable IP Network and MPLS VPNs

An MPLS network is typically implemented in a service provider or carrier network. It consists of interconnected routing devices known as Provider Edge (PE) routers and P (Provider) routers running MPLS services. The access networks, attached to the edge of the MPLS network via Customer Edge (CE) devices and PE routers, may be operated by regional Internet Service Providers (ISPs), local network operators, or even private companies.

- Attachment Circuit

In an MPLS Layer 2 VPN, the circuit or virtual circuit that links a CE device to a PE router. An Attachment Circuit may be a Frame Relay Data Link Circuit Identifier (DLCI), an ATM Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI), an Ethernet port, a Virtual Local Area Network (VLAN), or some other type of circuit or virtual circuit.

- BGP

Border Gateway Protocol. A routing protocol, defined in RFC 1657, that updates routes between autonomous systems.

- Binding

The process of associating an MPLS label with a Forwarding Equivalence Class (FEC). Control binding, which is a static form of binding, uses control messages (such as LDP) or specific predetermined commands and parameters to bind a label to an FEC.

- CE Device

Customer Edge Device. A device running in the customer's network (access network) that is connected to a service provider's PE router and is involved in an MPLS Layer 3 or Layer 2 VPN.

In an MPLS Layer 3 VPN, a CE is a router. In an MPLS Layer 2 VPN, a CE may be a router, a switch, a host, or any other device that the Layer 2 VPN customer needs to attach to the VPN.

- CR-LDP

Constraint-based Routing Label Distribution Protocol. An MPLS label signaling protocol used to advertise labels between PE and P routers to establish, maintain, and remove LSPs. CR-LDP is a revised version of LDP that includes traffic engineering extensions.

- FEC

Forwarding Equivalence Class. A group of IP packets that are forwarded through the MPLS network over the same path with the same priority and the same label; for example, all IP traffic going to the same subnet (say, 172.16). Each FEC defines a specific LSP and label.

An FEC can be based on a variety of access control list matches such as source address, destination address, BGP next hop, application type, and Differentiated Services (DiffServ) marking.

- Forwarder

connects a customer-side Attachment Circuit to an MPLS-side Pseudowire.

- to

- IBGP

Internal BGP. A session between two BGP peers in the same autonomous system, for the purpose of communicating externally derived routing information within the autonomous system. IBGP peers can be attached using a full-mesh topology or the Route Reflector model.

- Kompella VPN

An MPLS Layer 2 VPN based on the Kompella implementation methodologies. The signaling protocol used by Kompella-implemented VPNs is Multiprotocol BGP (MBGP).

- Label

A short identifier, attached to a packet, that identifies the path (LSP) that the labeled packet should take through the MPLS network. The label, a 20-bit unsigned integer in the range 0 through 1048575, is part of a 32-bit (4-byte) MPLS header that is inserted between the Layer 2 and Layer 3 headers of the packet by an ingress PE router.

A label contains an index into a forwarding table, which specifies the next hop for the packet. It is a shorthand notation that indexes the forwarding decision made by P routers.

- Label (or MPLS) Signaling Protocol

A signaling protocol between the PE/P routers to create, maintain, and delete LSPs. The protocol (LDP, CR-LDP, or RSVP-TE) is responsible for assigning labels, managing quality of service issues, and handling error conditions.

- **Label Stacking**
Adding multiple MPLS labels to a single packet. Label stacking is used for MPLS VPNs and when traversing multiple MPLS networks.
- **Label Swapping**
Using the incoming label to determine the outgoing label, encapsulation, and port; then replacing the incoming label with the outgoing label.

Label swapping takes place at P routers, not at ingress or egress PE routers. The swap operation consists of looking up the incoming label in the local label table to determine the outgoing label and the output port.
- **Label Table**
See MPLS Forwarding Table.
- **L2VPN**
See MPLS Layer 2 VPN.
- **L3VPN**
See MPLS Layer 3 VPN.
- **LDP**
Label Distribution Protocol. An MPLS label signaling protocol used to advertise labels between PE and P routers to establish, maintain, and remove LSPs. LDP is also used in Martini-implemented VPNs to exchange VPN reachability information between PE routers.
- **LER**
Label Edge Router. Essentially, an LER is a PE router without the software upgrade needed to support MPLS as a network-based VPN tunneling mechanism. See PE router.
- **LSP**
Label Switched Path. A concatenation of LSP hops that form an end-to-end forwarding path through the MPLS network. An LSP starts at an ingress PE router, crosses one or more P routers, and ends at an egress PE.

An LSP can be set up permanently by manually defining specific paths across a network for specific types of traffic, or set up on-the-fly using constraint-based routing based on parameters that constrain the forwarding direction. Constraint-based routing involves programming traffic-engineering parameters into the network.

- LSP Hop

See LSP Segment.

- LSP Segment

One hop between MPLS-enabled (PE/P) routers. An LSP consists of a set of defined hops between two PE routers. In the MPLS Manager environment, LSP incoming and outgoing segments represent incoming and outgoing labels in a PE/P router's MPLS forwarding table.

- LSR

Label Switching Router. An LSR is a P router. See P Router.

- Martini VPN

An MPLS Layer 2 VPN based on the Martini-implementation methodologies. The signaling protocol used by Martini-implemented VPNs is LDP.

- MBGP (also known as MP-BGP)

Multiprotocol Border Gateway Protocol. An extension to IBGP that allows the advertising of IPv6, multicast, and other non-IPv4 topologies within and between BGP autonomous systems. For MPLS Layer 3 VPNs and Kompella-implemented Layer 2 VPNs, MBGP is the mechanism used to distribute VPN-related information (such as VPN membership, reachability, topology, and tunnel endpoint information) between the PE routers.

- MBGP (or MP-BGP) Session

Multiprotocol Border Gateway Protocol Session. A link between PE routers in an MPLS network supporting MPLS Layer 3 VPNs or Kompella-implemented Layer 2 VPNs.

- MPLS

Multiprotocol Label Switching. A set of protocols developed by the Internet Engineering Task Force (IETF) that allows IP packets to be switched through the Internet by forwarding IP packets based on a short identifier known as a label. MPLS overcomes some of the shortcomings of IP networks through its ability to build virtual circuits called Label Switched Paths (LSPs) across IP networks. MPLS is also a key enabler for IP-based services such as Layer 3 VPNs.

Although originally designed to handle IP packets, MPLS can also handle non-IP packets via a Layer 2 VPN service by carrying subscriber Layer 2 frames from one customer site to another through LSPs and the MPLS backbone.

- MPLS FIB

MPLS Forwarding Information Base. See MPLS Forwarding Table.

- MPLS Forwarding Table

MPLS forwarding table, also known as the MPLS FIB or label table, is a label/interface look-up table used by PE routers to assign packets, received from CE routers/devices, to labels, and used by P routers to rapidly switch data traffic through the MPLS network.

- MPLS Layer 2 VPN

A provider-provisioned Layer 2 VPN, based on the Martini proposal, that supports MPLS as a network-based VPN tunneling mechanism at the Layer 2 level: Frame Relay, ATM, Ethernet, and so on. The MPLS Manager can discover both Martini-implemented and Kompella-implemented Layer 2 VPNs.

- MPLS Layer 3 VPN

A provider provisioned Layer 3 VPN, as defined by RFC-2547bis, that supports MPLS as a network-based VPN tunneling mechanism at the Layer 3 level. All functions associated with establishing, maintaining, and operating an MPLS VPN take place in the PE routers. Routing updates between PE routers are accomplished via MBGP.

- MPLS Network

MPLS network, also known as MPLS-enabled network or MPLS domain, is typically a large group of interconnected PE and P routers that span a large geographic area.

- **MPLS Service**

A router (PE, P) running MPLS software. MPLS service has a slightly different meaning in the MPLS Manager environment: MPLS Manager creates an MPLS service instance for each routing device discovered in the topology regardless of whether the device supports MPLS. The instance contains the device type: P, PE, CE, NON_MPLS, or Other.

- **NLRI**

Network Layer Reachability Information. The part of an MBGP routing update (control traffic) containing the VPN-IP address. For RFC 2547bis functionality, the NLRI represents a route to an arbitrary customer site or a set of customer sites within the VPN.

- **P Router**

Provider Router. An MPLS-capable router that participates in the establishment of LSPs based on pre-established IP routing information. It switches packets based on labels instead of making IP forwarding decisions. The incoming label instructs the P router where to forward the packets.

- **PE Router**

Provider Edge Router. An MPLS-capable router that operates at the edge of the access network and the MPLS network. It connects with one or more access networks (Frame Relay, ATM, Ethernet), determine routes, and adds or removes labels. It binds MPLS labels to FECs and LSPs, and handles and controls Layer 3 and Layer 2 VPN routing.

For an MPLS Layer 3 VPN, an ingress PE router examines the incoming packet's IP address, determines a route, assigns an LSP, and attaches two labels to the IP packet. The P routers in the MPLS network use the outer label to route the IP packet to the appropriate egress PE router. The egress PE router removes the outer label from the IP packet and forwards the IP packet to its destination via the inner label and standard IP routing.

For an MPLS Layer 2 VPN (and assuming a point-to-point VPN—see VPWS), an ingress PE router maps the incoming Layer 2 frame to an LSP and attaches two labels to the data frame. The P routers in the MPLS network use the outer label to route the Layer 2 frame to the appropriate egress PE router. The egress PE router removes the outer label from the Layer 2 frame and forwards the frame to its destination via the inner label.

The fact that two labels temporarily exist between the source and destination is completely transparent to the customer, the applications, and even the customer's networking equipment.

- Penultimate Hop Pop (PHP)

Penultimate hop pop, also known as penultimate label pop, is a process by which the penultimate router is directed to pop the outer label prior to forwarding the packet to the egress PE router. Using LDP, this action is accomplished by assigning the special label 3 (implicit Null label) as the outgoing label in the penultimate router's MPLS forwarding table.

- Penultimate Router

The last P router in an LSP. The penultimate router removes the outer label from a packet.

- Pseudowire

A bidirectional virtual connection that, in the MPLS environment, is carried over a pair of LSPs and terminated by a pair of Forwarders. A Pseudowire provides connectivity between two Attachment Circuits that are on the edges of the MPLS network.

- Route Distinguisher

An 8-byte value placed in front of a BGP IPv4 network route advertisement to identify the VRF to which the particular MPLS Layer 3 VPN route belongs. Typically, each VRF is assigned a unique route distinguisher, although it is common practice to assign the same route distinguisher to all the VRFs belonging to the same VPN. The route distinguisher is the means by which the PE router keeps track of overlapping customer IP address spaces.

A route distinguisher consists of a 2-byte Type field, a 2-byte Autonomous System Number (ASN) field, and a 4-byte Assigned Number field. Typically, only the ASN and Assigned Number fields are included in a route distinguisher; for example, 100:3000.

Note:

Because *route distinguisher* is included in MBGP routing updates, it is also relevant to Kompella-implemented Layer 2 VPNs.

- Route Target

A route target is essentially a VPN identifier. Route targets determine what routes a PE router exports from a VRF into BGP, and what routes a PE router imports from BGP into the VRF.

Each VRF has a list of route target communities with which it is associated; the list is defined for both export and import. The host PE router attaches the route target export list to each route advertised by the VRF. The host PE router adds a route to the VRF if the route target list attached to an advertised route contains at least one of the members in the VRF's route target import list.

The export list and import list implicitly determine the VPN topology. Implementing a simple VPN topology, such as full mesh, requires only one route target, whereas implementing a more complex VPN topology, such as hub and spoke, may require more than one route target. In the former case, a VRF's export list and import list contain the same route target. In the latter case, a VRF's export list and import list may contain different route targets.

Note:

Because *route target* is included in MBGP routing updates, it is also relevant to Kompella-implemented Layer 2 VPNs.

- RSVP-TE

Resource Reservation Protocol with Traffic Engineering extensions. An MPLS label signaling protocol used to advertise labels between PE and P routers to establish, maintain, and remove LSPs.

- TE

Traffic Engineering. The process of mapping traffic flows to paths other than the paths that would have been chosen by standard routing protocols. Traffic engineering can be achieved either manually or through a set of defined parameters whose requirements are then met by each appropriate network resource to establish the optimal path.

- Virtual Connection

A connection between end users that has a defined route and endpoints.

- VPN

Virtual Private Network. A private multi-site network created by using shared resources within a public network. No site outside the VPN can intercept packets or inject new packets into the VPN.

An MPLS Layer 3 VPN is a collection of VRFs that are members of the same VPN. An MPLS Layer 2 VPN is a collection of Forwarders and Pseudowires that are members of the same VPN.

- **VPN Path**
The data traffic path between two customer sites in a VPN.
- **VPN Peers**
A pair of peer VRFs hosted by different PEs and part of the same MPLS Layer 3 VPN.
- **VPN Site**
A VPN endpoint.
- **VPN Topology**
The way traffic is routed between the various sites within a VPN. Options include *full mesh* (where each customer site can communicate directly with every other customer site in the VPN), *hub and spoke* (where all traffic flows to/from a central hub site), and *partial mesh*. Essentially, a partial-mesh VPN is a hub-and-spoke VPN that has multiple hubs.
- **VPN-IP Address**
Virtual Private Network IP Address. An address consisting of an 8-byte route distinguisher and a 4-byte IPv4 address. A VPN-IP address identifies the VRF to which the particular VPN route belongs.
- **VPWS**
Virtual Private Wire Service. A point-to-point circuit (link) connecting two CE devices. A CE in the customer network is connected through an Attachment Circuit to a PE router in the MPLS network.
- **VRF**
VPN Routing and Forwarding. A forwarding table used by the PE routers to establish Layer 3 VPN paths through the MPLS network. A PE router maintains a separate VRF for each directly connected customer VPN site.

A VRF is configured with a name, a route distinguisher, a route target export list, and a route target import list. For example:

```
ip vrf CE
  rd 100:130
  route-target export 100:3000
  route-target import 100:3000
```

A VRF consists of an IP routing table, a derived forwarding table, a set of logical interfaces (tied to the locally attached customer VPN site) that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table.

RFC 2547bis specifies MBGP for intra-VRF route exchange. BGP updates are based on the export and import routing policies configured within each PE router.

- VRF Route Table

A table in a VRF that stores routing information for a particular VPN user. The table maps the VPN-IP route for the user to two labels: an outer label used to reach the PE router directly connected to the customer VPN site associated with the advertised NLRI, and an inner label used to reach the advertised NLRI.

B

Root-Cause Notifications from Availability Manager

The MPLS Manager receives root-cause notifications from the Availability Manager regarding failures in the underlying physical network. Table 47 lists these notifications.

Table 47: Notification to the MPLS Manager from the Availability Manager

ELEMENT	ROOT-CAUSE NOTIFICATION
Chassis	Down
Card	Down
Unitary Computer System (Router, Switch etc.)	Down Unstable
NetworkConnection	Down Unstable DownOrFlapping
Partition	Down
NetworkAdapter (Interface, Port)	Down Disabled

C

Polling for Analysis

The MPLS Manager uses Simple Network Management Protocol (SNMP) polling to obtain data for its correlation analysis. The parameters for controlling SNMP polling are accessed through the Polling and Thresholds Console.

SNMP Poller

The MPLS Manager uses a synchronous, multi-threaded SNMP poller. By default, the SNMP poller uses ten synchronous polling threads.

The SNMP poller fully supports the SNMP V1 and V2C protocols. With SNMP V1, the correlation model uses 32-bit counters in its correlation analysis. With SNMP V2C, the correlation model uses high-capacity 64-bit counters in its correlation analysis. Using 64-bit counters is critical for performance analysis of high-speed data links because using 32-bit counters might result in wrapping (overflow) of the counters between polls.

Note: The SNMP poller for MPLS Manager does not support SNMP V3.

Polling for devices with multiple IP addresses is supported because the SNMP poller supports multiple IP addresses for each SNMP agent. The SNMP poller automatically switches to an alternate IP address during failures, thereby ensuring the integrity of the SMARTS correlation analysis during an outage.

Just-In-Time Polling

The SNMP poller's MIB variable poll list is driven by a Just-In-Time polling algorithm, which ensures that only those MIB variables needed for correlation are polled. For example, if a port monitored for performance data is disabled, or goes down, the SNMP poller automatically removes the relevant MIB variables from the poll list. If the port is re-enabled, or comes back up, the variables are automatically put back onto the MIB poll list.

Request-Consolidation Polling

Issuing a single SNMP GET request that requests 10 variables is more efficient than issuing 10 GET requests each requesting a single variable. The SNMP poller consolidates as many variables as possible into a single SNMP GET request. The consolidation is not restricted to variables from the same SNMP table. Polling consolidation continually adapts to changes in the MIB variable poll list.

Upon encountering a non-fatal error during polling consolidation, the SNMP poller responds differently to an SNMP V1 agent than to an SNMP V2C for the following reason: Where an SNMP V1 agent *stops* processing a request upon encountering an error, an SNMP V2C agent *continues* processing a request upon encountering an error. An SNMP V2C agent handles errors on a per-OID basis.

If a non-fatal error is encountered by an SNMP V1 agent during a GET request seeking multiple variables, the SNMP poller suspends the polling of the affected variable because continuing to poll that variable would require the resending of the remainder of the request after receiving the error, which would probably impact the performance of the SNMP V1 agent; the SNMP poller continues to poll the unaffected variables. (An example of an affected variable is one that has become unavailable due to a configuration change.) This behavior enables the SNMP poller to operate efficiently with an SNMP V1 agent during unexpected changes to a device's configuration.

In contrast, if a non-fatal error is encountered by an SNMP V2C agent during a GET request seeking multiple variables, the SNMP poller continues the polling of the affected variable as well as the unaffected variables.

Wildcard Patterns

A wildcard pattern is a series of characters that are matched against incoming character strings. You can use these patterns when you define pattern matching criteria.

Matching is done strictly from left to right, one character or basic wildcard pattern at a time. Basic wildcard patterns are defined in Table 48. Characters that are not part of match constructs match themselves. The pattern and the incoming string must match completely. For example, the pattern *abcd* does not match the input *abcde* or *abc*.

A compound wildcard pattern consists of one or more basic wildcard patterns separated by ampersand (&) or tilde (~) characters. A compound wildcard pattern is matched by attempting to match each of its component basic wildcard patterns against the entire input string. For compound wildcard patterns, see Table 49.

If the first character of a compound wildcard pattern is an ampersand (&) or tilde (~) character, the compound is interpreted as if an asterisk (*) appeared at the beginning of the pattern. For example, the pattern *~*[0-9]** matches any string not containing any digits. A trailing instance of an ampersand character (&) can only match the empty string. A trailing instance of a tilde character (~) can be read as "except for the empty string."

Note: Spaces are interpreted as characters and are subject to matching even if they are adjacent to operators like "&".

Table 48: Basic Wildcard Patterns

CHARACTER	DESCRIPTION
Note: Spaces specified before or after wildcard operators are interpreted as characters and are subject to matching.	
?	Matches any single character. For example, <i>server?.smarts.com</i> matches <i>server3.smarts.com</i> and <i>serverB.smarts.com</i> , but not <i>server10.smarts.com</i> .
*	Matches an arbitrary string of characters. The string can be empty. For example, <i>server*.smarts.com</i> matches <i>server-ny.smarts.com</i> and <i>server.smarts.com</i> (an empty match).
[set]	Matches any single character that appears within [set]; or, if the first character of [set] is (^), any single character that is <i>not</i> in the set. A hyphen (-) within [set] indicates a range, so that [a-d] is equivalent to [abcd]. The character before the hyphen (-) must precede the character after it or the range will be empty. The character (^) in any position except the first, or a hyphen (-) at the first or last position, has no special meaning. Example, <i>server[789].smarts.com</i> matches <i>server7.smarts.com</i> through <i>server9.smarts.com</i> , but not <i>server6.smarts.com</i> . It also matches <i>server-.smarts.com</i> . Example: <i>server[^12].smarts.com</i> does not match <i>server1.smarts.com</i> or <i>server2.smarts.com</i> , but will match <i>server8.smarts.com</i> .
<n1-n2>	Matches numbers in a given range. Both <i>n1</i> and <i>n2</i> must be strings of digits, which represent non-negative integer values. The matching characters are a non-empty string of digits whose value, as a non-negative integer, is greater than or equal to <i>n1</i> and less than or equal to <i>n2</i> . If either end of the range is omitted, no limitation is placed on the accepted number. For example, <i>98.49.<1-100>.10</i> matches a range of IP addresses from <i>98.49.1.10</i> through <i>98.49.100.10</i> . Example of an omitted high end of the range: <i><50></i> matches any string of digits with a value greater than or equal to 50. Example of an omitted low end of the range: <i><-150></i> matches any value between zero and 150. A more subtle example: The pattern <i><1-10>* </i> matches 1, 2, up through 10, with <i>*</i> matching no characters. Similarly, it matches strings like <i>9x</i> , with <i>*</i> matching the trailing <i>x</i> . However, it does not match <i>11</i> , because <i><1-10></i> always extracts the longest possible string of digits (11) and then matches only if the number it represents is in range.
	Matches alternatives. For example, <i>"ab bc cd"</i> without spaces matches exactly the three following strings: <i>"ab"</i> , <i>"bc"</i> , and <i>"cd"</i> . A as the first or last character of a pattern accepts an empty string as a match. Example with spaces <i>"ab bc"</i> matches the strings <i>"ab "</i> and <i>" bc"</i> .
\	Removes the special status, if any, of the following character. Backslash (\) has no special meaning within a set ([set]) or range (<n1-n2>) construct.

Special characters for compound wildcard patterns are summarized below.

Table 49: Compound Wildcard Patterns

CHARACTER	DESCRIPTION
&	<p>"And Also" for a compound wildcard pattern. If a component basic wildcard pattern is preceded by & (or is the first basic wildcard pattern in the compound wildcard pattern), it <i>must</i> successfully match.</p> <p>Example: *NY*&*Router* matches all strings which contain NY and also contain Router.</p> <p>Example: <1-100>&*[02468] matches even numbers between 1 and 100 inclusive. The <1-100> component only passes numbers in the correct range and the *[02468] component only passes numbers that end in an even digit.</p> <p>Example: *A* *B*&*C* matches strings that contain either an A or a B, and also contain a C.</p>
~	<p>"Except" for a compound wildcard pattern (opposite function of &). If a component basic wildcard pattern is preceded by ~, it <i>must not</i> match.</p> <p>Example: 10.20.30.*~10.20.30.50 matches all devices on network 10.20.30 except 10.20.30.50.</p> <p>Example: *Router*~*Cisco*&*10.20.30.*~10.20.30.<10-20>* matches a Router, except a Cisco router, with an address on network 10.20.30, except not 10.20.30.10 through 10.20.30.20.</p>

Index

A

- Adapter
 - Cisco ISP 33
- Adding or removing a setting 91
- Adding or removing matching criteria 92
- Analysis 8
- AnalysisMode
 - CE to CE Ping setting 76
 - Forwarder SNMP setting 73
 - LDP Session SNMP setting 74
 - PE to CE Ping setting 78
 - PE to PE Ping setting 80
 - PE to Unmanaged CE Ping setting 82
 - PE to VRF Ping setting 84
 - VRF SNMP setting 75
- Attachment Circuit 4
- Attributes
 - LSP 15
 - LSPHop 16
 - LSPInSegment 17
 - LSPOutSegment 18
 - MPLSService 13
 - RouteTarget 24
 - VPN 20, 26
 - VRF 22
- Availability Manager 1, 2

B

- BASEDIR xii

C

- CE to CE Ping setting
 - AnalysisMode 76
 - Classification 76
 - DestinationIP 76
 - DestinationName 76
 - packetSize 76
 - ping_TimeOut 77
 - polling_TimeOut 77
 - pollingInterval 77
 - SourceName 77
 - SourceType 77

- Changing matching criteria 93
- Changing priority of a group 92
- Changing setting parameters 93
- Cisco ISC adapter 33
- Classification
 - CE to CE Ping setting 76
 - PE to CE Ping setting 78
 - PE to PE Ping setting 80
 - PE to Unmanaged CE Ping setting 82
 - PE to VRF Ping setting 84
- Containment 65
- Containment dialog 66
 - Opening 66
- Copying a group 94
- Creating a group 95

D

- DestinationIP
 - CE to CE Ping setting 76
 - PE to CE Ping setting 78
 - PE to PE Ping setting 80
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
- DestinationName
 - CE to CE Ping setting 76
 - PE to CE Ping setting 79
 - PE to PE Ping setting 81
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
- DeviceType attribute 13
- Disconnected events
 - LdpAdjacency
 - Disconnected 12, 32
 - PseudoWire
 - Disconnected 12, 30
- Discovery 2
 - MPLS Manager 3
- Domain Manager
 - Containment 65
- Domain Manager Administration Console 87
- Down events
 - LdpAdjacency
 - Down 12, 32

PseudoWire
Down 12, 30, 41

E

Edges
Table of 53
Element
Assigning to groups 90

F

Forwarder 26
Impacted 40
IsDown 12, 28, 41
Forwarder SNMP setting
AnalysisMode 73
PollingInterval 73
Retries 73
Timeout 73

G

Global Console 8, 49
Containment dialog 66
Domain Manager Administration Console 87
Map Console view 50, 53
Notification Log Console view 50
Notification Properties dialog 50, 67
Polling and Thresholds Console 88

Global Manager 1, 8

Group

Assigning members 90
Changing priority 92
Copying 94
Creating 95
Definition of 70
Properties 90

Groups

Polling
Forwarders 71
LdpProtocolEndpoint 71
VRFs 71
Remote ping
CE to CE Pings 71
PE to CE Pings 71
PE to PE Pings 71
PE to Unmanaged CE Pings 71
PE to VRF Pings 72
System Write Community Strings 72

I

ICMP Echo requests 7
ifIndex attribute 18
Impact event
Forwarder Impacted 40
LdpProtocolEndpoint Impacted 40
LSP 41
Notification 40
PseudoWire Impacted 40
Summary table 40, 41, 104
VLAN Impacted 41
VPN Impacted 40
VRF Impacted 40
Infrastructure
Containment 65
IsDown events
Forwarder
IsDown 12, 28, 41
LdpProtocolEndpoint
IsDown 12, 31

K

Key attribute 18, 24

L

Label attribute 16, 17, 18
Label Switched Path (LSP) 14
LDP Session SNMP setting
AnalysisMode 74
PollingInterval 74
Retries 74
Timeout 74
LDPAdjacency 31
LdpAdjacency 5
Disconnected 12, 32
Down 12, 32
LdpProtocolEndpoint 5, 30
Impacted 40
IsDown 12, 31
LSP 14
Attributes 15
LSPId 15
Name 15
Impacted 41
LSP Hops Map 58
LSP Map 56
LSPHop 15
Attributes 16
Label 16

- LSPId 16
 - Name 16
- LSPId attribute 15, 16, 17, 18
- LSPInSegment 17
 - Attributes 17
 - Label 17
 - LSPId 17
 - Name 17
- LSPOutSegment 17
 - Attributes 18
 - ifIndex 18
 - Key 18
 - Label 18
 - LSPId 18
 - Name 19
 - NextHopIP 19

M

- Map Console
 - Icons and indicators 53
 - Type of map 55
- Maps
 - LSP 56
 - LSP Hops 58
 - MPLS topology 53
 - Opening 53
 - Type of 55
 - VPN 59, 61, 64
- Matching
 - Pattern 129
- Matching criteria
 - Adding or removing 92
 - Changing 93
- MaxRoutes
 - VRF External setting 75
- MaxRoutes attribute 22
- MidRoute Threshold
 - VRF External setting 75
- MidRouteThreshold attribute 22
- Misconfiguration events
 - RouteTarget
 - Misconfiguration 11, 24
 - Summary table 11
 - VRF
 - Down 11, 23
 - MaxRoutesReached 11, 23
 - NoRoutes 11, 23
 - WarningThresholdCrossed 11, 23
- Monitoring 8

- MPLS Domain 3
- MPLSService 13
 - Attributes 13
 - DeviceType 13
 - Name 13
 - NumberOfVRFs 14
 - Supports_LSR_MIB 14
 - Supports_VPN_MIB 14
 - TotalVRFRoutes 14

N

- Name attribute 13, 15, 16, 17, 19, 20, 22, 24, 26
- NextHopIP attribute 19
- Nodes
 - Table of 53
- Notification Log Console 8
- Notification Properties dialog 50, 67
 - Opening 50
- Notifications
 - Impact events 40
 - Misconfiguration events 11
- NumberOfRoutes attribute 22
- NumberOfVRFs attribute 14

O

- Opening a Containment dialog 66
- Opening a Map 53
- Opening a Notification Properties dialog 50
- Operator
 - Wildcard 130

P

- packetSize
 - CE to CE Ping setting 76
 - PE to CE Ping setting 79
 - PE to PE Ping setting 81
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
- Pattern 129
- Pattern matching 129
- PE to CE Ping setting
 - AnalysisMode 78
 - Classification 78
 - DestinationIP 78
 - DestinationName 79
 - packetSize 79
 - ping_TimeOut 79
 - polling_TimeOut 79
 - pollingInterval 79

- SourceName 79
 - SourceType 79
 - PE to PE Ping setting
 - AnalysisMode 80
 - Classification 80
 - DestinationIP 80
 - DestinationName 81
 - packetSize 81
 - ping_TimeOut 81
 - polling_TimeOut 81
 - pollingInterval 81
 - SourceName 81
 - SourceType 81
 - PE to Unmanaged CE Ping setting
 - AnalysisMode 82
 - Classification 82
 - DestinationIP 83
 - DestinationName 83
 - packetSize 83
 - ping_TimeOut 83
 - polling_TimeOut 83
 - pollingInterval 83
 - SourceName 83
 - SourceType 83
 - PE to VRF Ping setting
 - AnalysisMode 84
 - Classification 84
 - DestinationIP 85
 - DestinationName 85
 - packetSize 85
 - ping_TimeOut 85
 - polling_TimeOut 85
 - pollingInterval 85
 - SourceName 85
 - SourceType 85
 - Penultimate hop popping 15
 - ping_TimeOut
 - CE to CE Ping setting 77
 - PE to CE Ping setting 79
 - PE to PE Ping setting 81
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
 - Polling
 - Groups 72, 73, 74
 - SNMP 8, 69, 128
 - Polling and Thresholds Console 87
 - Layout 88
 - Polling tab 88
 - Thresholds tab 88
 - Toolbar buttons 89
 - Polling tab 88
 - polling_TimeOut
 - CE to CE Ping setting 77
 - PE to CE Ping setting 79
 - PE to PE Ping setting 81
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
 - PollingInterval
 - Forwarder SNMP setting 73
 - LDP Session SNMP setting 74
 - VRF SNMP setting 75
 - pollingInterval
 - CE to CE Ping setting 77
 - PE to CE Ping setting 79
 - PE to PE Ping setting 81
 - PE to Unmanaged CE Ping setting 83
 - PE to VRF Ping setting 85
 - Priority
 - Changing 92
 - Provisioning system adapter 33
 - PseudoWire
 - Disconnected 12, 30
 - Down 12, 30, 41
 - Impacted 40
 - Pseudowire 4, 28
- ## R
- Remote Ping 7
 - Remote ping
 - Groups 76, 78, 80, 82, 84, 86
 - SNMP 69
 - REMOTEPING.conf file 87
 - Removing or adding a setting 91
 - Removing or adding matching criteria 92
 - Restoring default values of a setting 94
 - Retries
 - Forwarder SNMP setting 73
 - LDP Session SNMP setting 74
 - VRF SNMP setting 75
 - RouteDistinguisher attribute 22
 - Router Community Strings setting
 - WriteCommunity 86
 - RouteTarget 24
 - Attributes 24
 - Key 24
 - Name 24
 - Misconfiguration 11, 24

S

serverConnect.conf 87

Setting

Adding or removing 91

Changing parameters 93

Definition of 70

Restoring default values 94

SNMP

Polling 8, 69, 128

Remote ping 69

SourceName

CE to CE Ping setting 77

PE to CE Ping setting 79

PE to PE Ping setting 81

PE to Unmanaged CE Ping setting 83

PE to VRF Ping setting 85

SourceType

CE to CE Ping setting 77

PE to CE Ping setting 79

PE to PE Ping setting 81

PE to Unmanaged CE Ping setting 83

PE to VRF Ping setting 85

Supports_LSR_MIB attribute 14

Supports_VPN_MIB attribute 14

T

Thresholds tab 88

Timeout

Forwarder SNMP setting 73

LDP Session SNMP setting 74

VRF SNMP setting 75

Topology

Containment 65

Topology attribute 21, 26

TotalVRFRoutes attribute 14

Transparent LAN Service (TLS) 6, 46

Tunnel 28

V

Virtual Private Network (VPN) 20, 25

VLAN

Impacted 41

VPN 20, 25

Attributes 20, 26

Name 20, 26

Topology 21, 26

VPNTYPE 21, 26

Impacted 40

VPN Domain 4

VPN Map 59, 61, 64

VPN Routing and Forwarding (VRF) 21

VPNTYPE attribute 21, 26

VRF 21

Attributes 22

MaxRoutes 22

MidRouteThreshold 22

Name 22

NumberOfRoutes 22

RouteDistinguisher 22

VRFName 22

Down 11, 23

Impacted 40

MaxRoutesReached 11, 23

NoRoutes 11, 23

WarningThresholdCrossed 11, 23

VRF External setting

MaxRoutes 75

MidRoute Threshold 75

VRF polling

VRF External setting 75

VRF SNMP setting 74

VRF SNMP setting

AnalysisMode 75

PollingInterval 75

Retries 75

Timeout 75

VRFName attribute 22

W

Wildcard 129

Chart of operators 130

WriteCommunity

Router Community Strings setting 86

