



MPLS Manager

1.1

USER'S GUIDE

P/N 300-002-533

REV A01

EMC Smarts

Corporate Headquarters:

Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright 1996-2005 by EMC Corporation ("EMC"). All rights reserved.

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

The Software and all intellectual property rights related thereto constitute trade secrets and proprietary data of EMC and any third party from whom EMC has received marketing rights, and nothing herein shall be construed to convey any title or ownership rights to you. Your right to copy the software and this documentation is limited by law. Making unauthorized copies, adaptations, or compilation works is prohibited and constitutes a punishable violation of the law. Use of the software is governed by its accompanying license agreement.

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 InCharge products mentioned in this publication are covered by one or more of the following U.S. Patent Nos. or pending patent applications: 5,528,516, 5,661,668, 6,249,755, 6,868,367 and 11/034,192.

"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. 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 portions of the software are set forth in the Third-Party Copyright Read Me file.

Contents

Preface	vii
Intended Audience	vii
Prerequisites	vii
Document Organization	viii
Documentation Conventions	ix
InCharge MPLS Management Suite Installation Directory	x
MPLS Management Suite Products	x
Additional Resources	x
InCharge Commands	x
Documentation	xi
Technical Support	xii
1 Introduction	1
Discovery	2
Availability Manager Discovery	2
MPLS Manager Discovery	2
Monitoring	4
Analysis	4
Notifications	4
2 MPLS and VPN Elements and Their Failures	5
MPLS and VPN Elements and Relationships	6
Summary of Misconfiguration Events	7
MPLSService	8
Attributes for MPLSService	8
LSP	9
Attributes for LSP	10
LSPHop	10

Attributes for LSPHop	11
LSPInSegment	12
Attributes for LSPInSegment	12
LSPOutSegment	13
Attributes for LSPOutSegment	13
VPN	14
Attributes for VPN	15
VRF	15
Attributes for VRF	16
Misconfiguration Events for VRF	17
RouteTarget	18
Attributes for RouteTarget	19
Misconfiguration Events for RouteTarget	19
Provisioning System Adapter Events	19
Underlying Transport Network Failures	20
3 MPLS Cross-Domain Impact Correlation Analysis	21
Impact Analysis Overview	21
Impact Analysis Model	22
Impact Analysis Events	24
Impact Analysis Examples	24
Example 1: CE-PE Router NetworkConnection Down	25
Example 2: PE Router Down	27
4 Viewing MPLS Notifications and Maps	29
Viewing MPLS Notifications	29
Opening an MPLS Notification Properties Dialog Box	30
MPLS Notification Properties	30
Viewing MPLS Topology in Maps	31
Opening an MPLS Topology Map	32
MPLS Topology Map Graphical Representations	32
MPLS Map Types	34

5	Customizing MPLS Polling	39
	Polling Groups and Settings	39
	Polling Groups	40
	Polling Settings	40
	Threshold Groups and Settings	41
	Opening the Polling and Thresholds Console	41
	Layout of the Polling and Thresholds Console	42
	Polling and Thresholds Console Toolbar Buttons	43
	Working With Polling Groups and Settings	44
	How Managed Elements Are Assigned to Groups	44
	Modifying the Properties of a Group	44
	Adding or Removing Group Settings	45
	Modifying the Priority of Groups	45
	Editing Matching Criteria	46
	Modifying the Parameters of a Setting	47
	Creating New Polling Groups	48
A	MPLS Terminology	49
B	MIBs Polled	59
	MIB Support for Cisco Devices and Juniper M/T Devices	59
	MIB Support for Juniper ERX Devices	60
C	CLI Commands	61
	CLI Support for Cisco Devices and Juniper M/T Devices	61
	CLI Support for Juniper ERX Devices	62
D	Polling for Analysis	63
	SNMP Poller	63
	Just-In-Time Polling	64
	Request-Consolidation Polling	64

E Wildcards	65
--------------------	-----------

Index	69
--------------	-----------

Preface

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

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 SMARTS software must be installed:

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

For information about installing these products, see the *InCharge IP Management Suite Installation Guide*, the *InCharge Service Assurance Management Suite Installation Guide*, and the *InCharge 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 MPLS Manager.
2. MPLS AND VPN ELEMENTS AND THEIR FAILURES	Describes the MPLS and VPN elements managed by MPLS Manager and identifies the root-cause problems and symptomatic events for each element type.
3. MPLS CROSS-DOMAIN IMPACT CORRELATION ANALYSIS	Describes how 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 AND MAPS	Describes how to use the Global Console to view MPLS notifications and topology maps for MPLS Manager.
5. CUSTOMIZING MPLS POLLING	Provides information and procedures for customizing SNMP polling for MPLS Manager.
A. MPLS TERMINOLOGY	Describes basic terms and concepts for MPLS and MPLS VPNs.
B. MIBS POLLED	Identifies the MIBs used by MPLS Manager to discover and monitor MPLS and VPN elements.
C. CLI COMMANDS	Identifies the CLI commands used by MPLS Manager to discover MPLS and VPN elements.
D. POLLING FOR ANALYSIS	Describes the SNMP polling engine used by MPLS Manager for correlation analysis.
E. WILDCARDS	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 SMARTS programs such as Domain Managers, Global Managers, and adapters.

InCharge MPLS Management Suite Installation Directory

In this document, the term **BASEDIR** represents the location where 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 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 SMARTS software, refer to the *InCharge System Administration Guide*.

MPLS Management Suite Products

The MPLS Management Suite offers the following products:

- MPLS Manager
- InCharge Adapter for Cisco ISC
- Perl API

Additional Resources

In addition to this document, SMARTS provides the following resources.

InCharge Commands

Descriptions of SMARTS commands 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 SMARTS documentation (also available in the **BASEDIR**/*smarts/doc/pdf* directory) helpful.

SMARTS Documentation

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

- *InCharge Release Notes*
- *InCharge Documentation Roadmap*
- *InCharge System Administration Guide*
- *InCharge ICIM Reference*
- *InCharge Dynamic Modeling Tutorial*
- *InCharge MODEL Reference Guide*
- *InCharge ASL Reference Guide*
- *InCharge Perl Reference Guide*

MPLS Management Suite Documentation

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

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

Refer to the *InCharge Documentation Roadmap* for documentation resources provided with other 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 EMC's secure extranet for customers and partners. Powerlink is an essential tool for obtaining web-based support from EMC. 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 Virtual Private Networks (VPNs).

The MPLS Manager enables service providers to fully realize the many benefits of MPLS, and provides them with the ability to intelligently manage and ensure the reliability of their core MPLS-based service offerings.

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

- Discovers, models, and monitors the network, MPLS, and VPN elements in the managed MPLS environment
- Identifies common configuration errors that occur when deploying and maintaining MPLS and VPN networks
- Correlates underlying network problems with MPLS and VPN impairments
- Reports its analysis results to the Global Manager

The MPLS Manager architecture is illustrated and described in the *InCharge 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 to the customer site routers (CE routers, Figure 1). It 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.

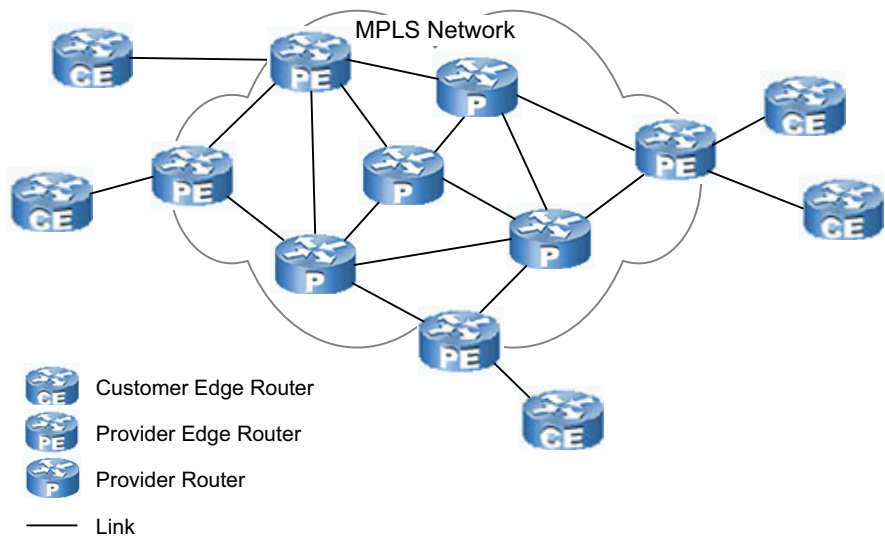


Figure 1: The Transport Domain—Discovered by the Availability Manager

MPLS Manager Discovery

The MPLS Manager discovers the MPLS logical topology (Figure 2) and the VPN logical topology (Figure 3) and models that topology in its repository. It 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

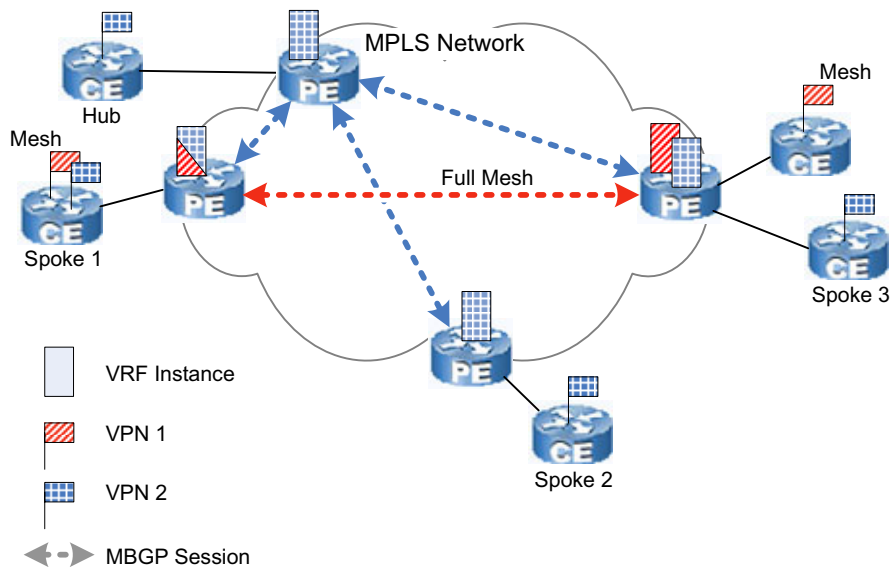


Figure 3: The VPN Domain—Discovered by the MPLS Manager

Monitoring

After the initial discovery of the MPLS and VPN elements, the MPLS Manager continuously monitors the status of the VPN Routing and Forwarding (VRF) elements, as well as the number of VRF routes maintained by the VRF elements, by periodically sending SNMP polls to the Provider Edge (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 MPLS Polling](#) on page 39.

Analysis

The MPLS Manager detects common configuration 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, 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 the misconfiguration 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 misconfiguration 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 5 and [MPLS Cross-Domain Impact Correlation Analysis](#) on page 21. For information about viewing the notifications, see [Viewing MPLS Notifications and Maps](#) on page 29.

2

MPLS and VPN Elements and Their Failures

This chapter describes the MPLS and VPN elements discovered and managed by the MPLS Manager and the misconfiguration events notified for the elements. In addition, it includes descriptions of MPLS and VPN element attributes.

The MPLS Manager identifies common configuration errors that occur when deploying and maintaining MPLS and VPN networks, and reports the errors to the Global Manager. It also correlates MPLS and VPN impairments with transport root-cause failures received from the Availability Manager to identify MPLS and VPN impacts, as explained in [MPLS Cross-Domain Impact Correlation Analysis](#) on page 21, 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 and VPN elements in the managed MPLS environment. This model represents the MPLS elements, their relationships, and their connections.

The elements in the discovered MPLS and VPN topology are represented as instances of the following InCharge Common Information Model (ICIM) classes:

- MPLSService
- LSP
- LSPHop
- LSPInSegment
- LSPOutSegment
- VPN
- VRF
- RouteTarget

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 *InCharge IP Discovery Guide Supplement for MPLS*. For general descriptions of ICIM classes, relationships, and connections, see the *InCharge ICIM Reference*.

Summary of Misconfiguration Events

MPLS Manager creates a misconfiguration event notification for each configuration error that it detects. Notifications are displayed in the Global Console.

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

Table 3: Misconfiguration Events Notified by MPLS Manager

MANAGED ELEMENT	EVENT	CONDITION
VRF	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: 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.		

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 49.

The relationships created for an MPLS service depend on the type of device (PE, P, CE) hosting the MPLS service. Consider the following three relationships created for MPLS service:

- **AttachedTo**—Applicable to an MPLS service hosted by a CE router. This relationship points to the VRF and PE router to which the CE router is attached.
- **MPLSInterfaces**—Applicable to an MPLS service hosted by a PE or P router. This relationship points to all the interfaces on the PE or P router that are MPLS-enabled.
- **VRFInterfaces**—Applicable to an MPLS service hosted by a PE router. This relationship points to all the interfaces on the PE router that are associated with VRFs.

Attributes for MPLSService

Table 4 lists key attributes for MPLSService.

Table 4: Attributes for MPLSService

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
DeviceType	Type of router hosting this MPLS service.	Enum: <ul style="list-style-type: none"> • PE • P • CE • NON_MPLS • Other
Name	Name assigned to this MPLS service. The name format is <code>MPLS-<i><dev></i></code> , where <i><dev></i> is the name or IP address of the router hosting this MPLS service.	String
NumberOfVRFs	Number of VRFs maintained by the router hosting this MPLS service. Applicable to PE routers only; otherwise, set to 0.	Integer

Table 4: Attributes for MPLSService (continued)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
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.

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

Attributes for LSP

Table 5 lists key attributes for LSP.

Table 5: 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. Typically, the IP address is 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 49. 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 6 lists key attributes for LSPHop.

Table 6: Attributes for LSPHop

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Label	<p>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.</p> <p>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.</p>	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	<p>Name assigned to this LSP hop. The name format is LSPHop- <label>/<destdev>/<destif>, where:</p> <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. <p>For an LSP hop having Label = 3, the name format is LSPHop-POP/<srcdev>/<srcif>-<key>, where:</p> <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 MPLS Manager: An LSP starts at, and is determined by, the LSP outsegment chosen by the ingress PE router.

Attributes for LSPInSegment

Table 7 lists key attributes for LSPInSegment.

Table 7: 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 8 lists key attributes for LSPOutSegment.

Table 8: 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 8: Attributes for LSPOutSegment (continued)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Name	<p>Name assigned to this LSP outsegment. The name format is LSPOutSegment-<i><label>/<dev>/<ifindex></i>, where:</p> <ul style="list-style-type: none"> • <i><label></i> is the label assigned to this LSP outsegment. • <i><dev></i> is the name or IP address of the router hosting the MPLS forwarding table to which this LSP outsegment belongs. • <i><ifindex></i> is the ifIndex attribute value for this LSP outsegment. <p>For an LSP outsegment having Label = 3, the name format is LSPOutSegment-POP/<i><dev>/<ifindex>-<key></i>, where:</p> <ul style="list-style-type: none"> • <i><dev></i> is the name or IP address of the router hosting the MPLS forwarding table to which this LSP outsegment belongs. • <i><ifindex></i> is the ifIndex attribute value for this LSP outsegment. • <i><key></i> 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

VPN

A VPN 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 MPLS Layer 3 VPN take place in the PE routers.

The P routers are not aware of the VPNs; they forward packets over the established LSPs. Similarly, the CE routers are not aware of the VPNs; 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.

Attributes for VPN

Table 9 lists key attributes for VPN.

Table 9: Attributes for VPN

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
Name	Name assigned to this VPN. For a full-mesh VPN, the name format is <code>VPN-<rt1></code> , where <code><rt1></code> is the name of the route target associated with this VPN. For a hub-and-spoke VPN, the name format is <code>VPN-<rt2>:<rt3></code> , where: <ul style="list-style-type: none"> • <code><rt2></code> is the name of the one route target associated with this VPN. • <code><rt3></code> is the name of the other route target associated with this VPN. 	String
Topology	Topology of this VPN, which is determined by examining the route targets exported and imported by the VRFs comprising this VPN.	String: <ul style="list-style-type: none"> • FullMesh • Hub&Spoke

VRF

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

Table 10: 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.

Table 10: Attributes for VRF (*continued*)

ATTRIBUTE	DESCRIPTION	ALLOWED VALUES
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

Misconfiguration Events for VRF

Table 11 lists the misconfiguration events notified for VRF.

Table 11: Misconfiguration Events for VRF

SYMPTOM	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.
* MPLS Manager also notifies VRF impacts for each of the down interfaces.	

Table 11: Misconfiguration Events for VRF (*continued*)

SYMPTOM	DESCRIPTION
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 or exceeds 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

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 12 lists key attributes for RouteTarget.

Table 12: 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

Misconfiguration Events for RouteTarget

A single misconfiguration 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.

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: InCharge Adapter for Cisco ISC (Cisco ISC Adapter), which interfaces with the Cisco Internet Solutions Center (ISC) provisioning system.

When the Cisco 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 *InCharge MPLS Manager User's Guide for Cisco ISC Adapter*.

Underlying Transport Network Failures

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

To understand how MPLS Manager correlates the underlying transport network failures received from Availability Manager with MPLS symptoms to identify MPLS impacts, see [MPLS Cross-Domain Impact Correlation Analysis](#) on page 21.

3

MPLS Cross-Domain Impact Correlation Analysis

The MPLS Manager calculates the impacts on MPLS and VPN 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 Availability Manager, including Router Down, Interface Down, Interface Disabled, NetworkConnection Down, and others. When the MPLS Manager receives a problem event, it correlates the problem with VPN, VRF, 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 Availability Manager.

Figure 4 shows the flow of information between the components in an MPLS Manager deployment to achieve MPLS and global impact analysis.

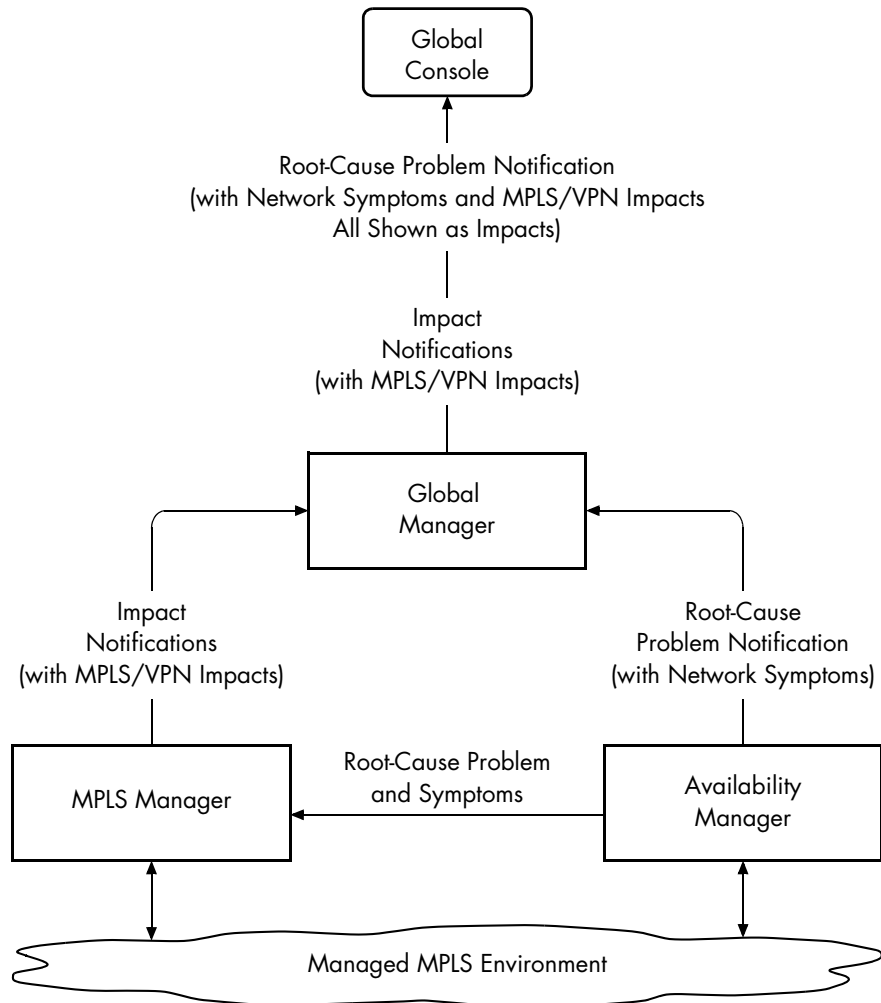


Figure 4: MPLS and Global Impact Analysis

Impact Analysis Model

Figure 5 demonstrates at a high level how MPLS Manager models (represents) the managed MPLS environment. The underlying transport network elements in the model, shown as white text on black background, are managed by Availability Manager. As such, MPLS Manager receives the status of these elements from Availability Manager.

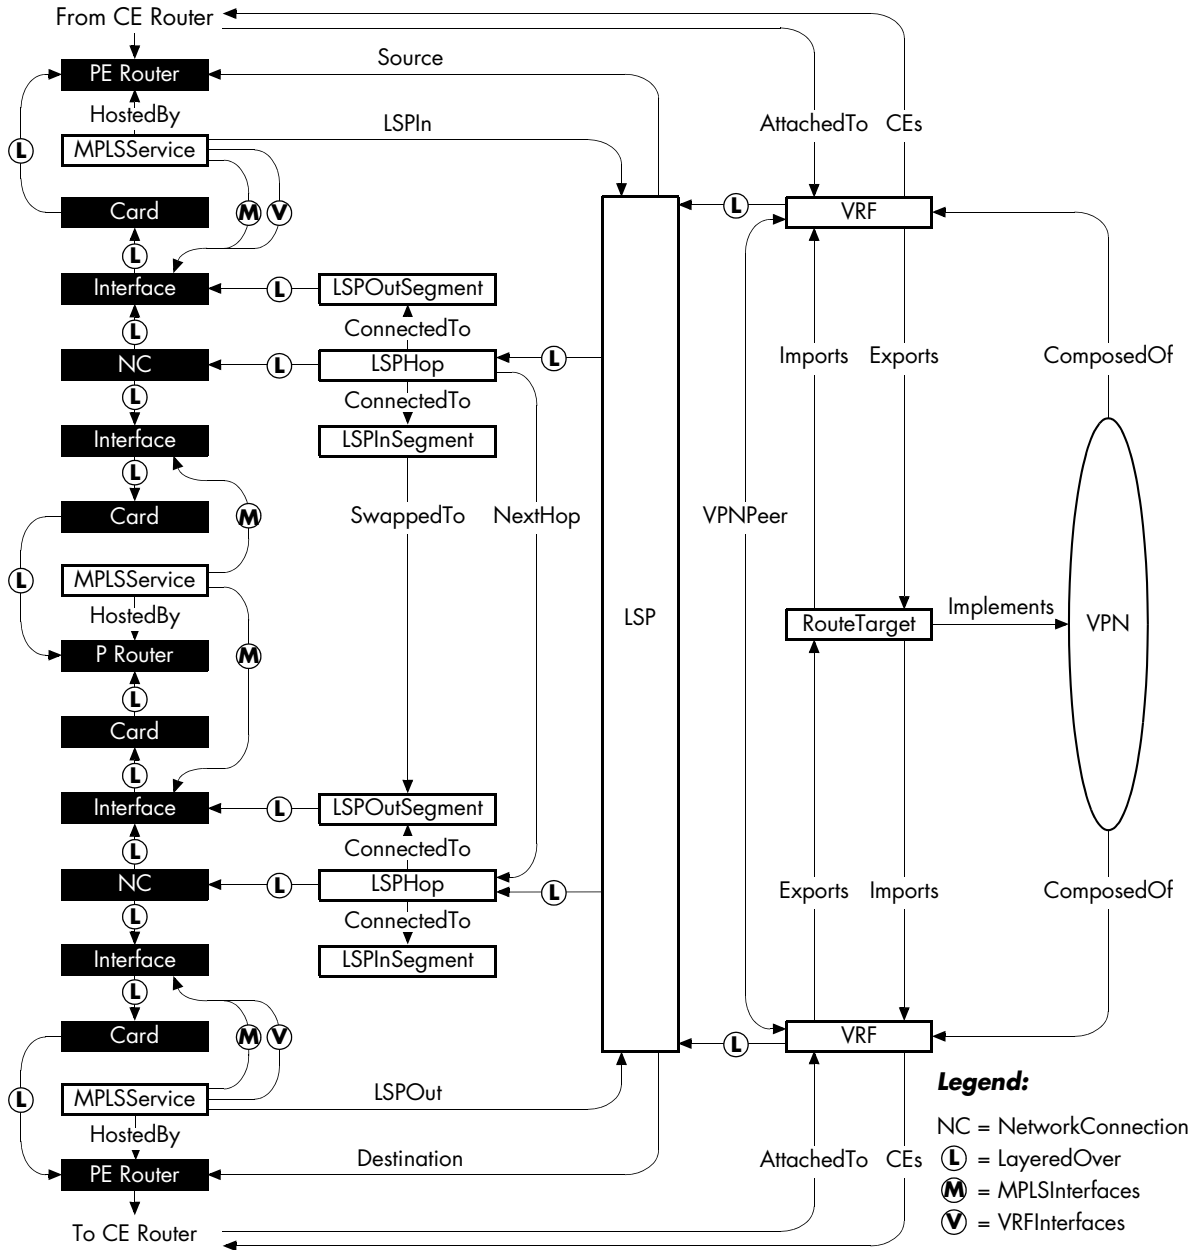


Figure 5: MPLS Manager View of its Managed Topology

Impact Analysis Events

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

Table 13 lists the impact events notified by MPLS Manager, including the condition for each event. The table also identifies the managed elements for which the impact events are notified.

Table 13: Impact Events Notified by MPLS Manager

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

The two examples that follow show how 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 MPLS Manager leaves off to perform global impact analysis.

Figure 6 shows the underlying transport failures for the two examples.

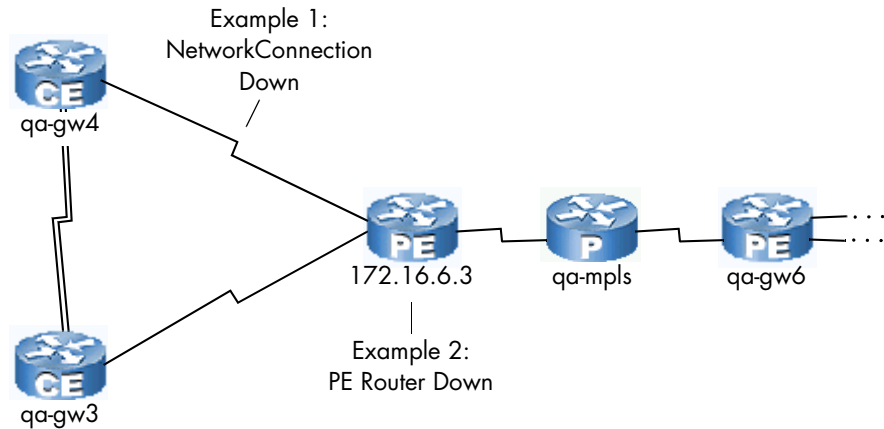


Figure 6: Underlying Transport Failures for the Examples

Example 1: CE-PE Router NetworkConnection Down

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

When Availability Manager detects the NetworkConnection Down problem, it sends the problem along with the network symptoms to both the Global Manager and MPLS Manager. 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 MPLS Manager to the symptoms in the root-cause problem notification received from 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 7) 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

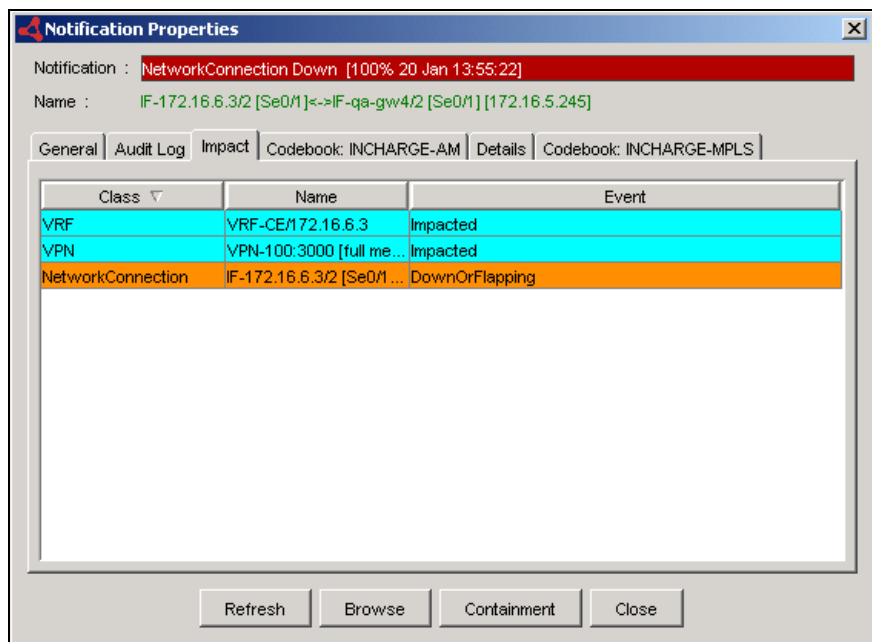


Figure 7: Notification Properties Dialog Box Showing a NetworkConnection Down Problem

For instructions on viewing detailed notification information, see [Viewing MPLS Notifications and Maps](#) on page 29.

Example 2: PE Router Down

In this example, MPLS Manager receives a Router Down problem from 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 8) 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

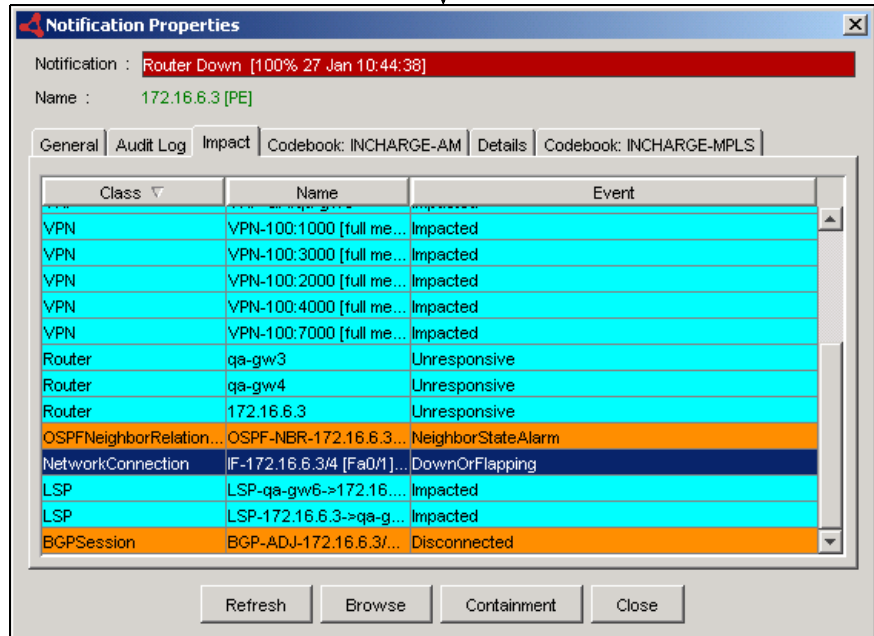


Figure 8: Notification Properties Dialog Box Showing a Router Down Problem

4

Viewing MPLS Notifications and Maps

This chapter describes using the Global Console to view MPLS notifications and the following MPLS topology maps:

- LSP maps
- LSP Hops maps
- VPN maps

You view notifications and topology maps by attaching the Global Console to the Global Manager.

For instructions on viewing notifications and maps, see the *InCharge 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 5.

Viewing MPLS Notifications

MPLS Manager reports notifications to the Global Manager, and the Global Manager combines these notifications with the notifications received from 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 9 and Figure 10 provide examples of detailed VRF Down and VRF NoRoutes information displayed in the Notification Properties dialog box.

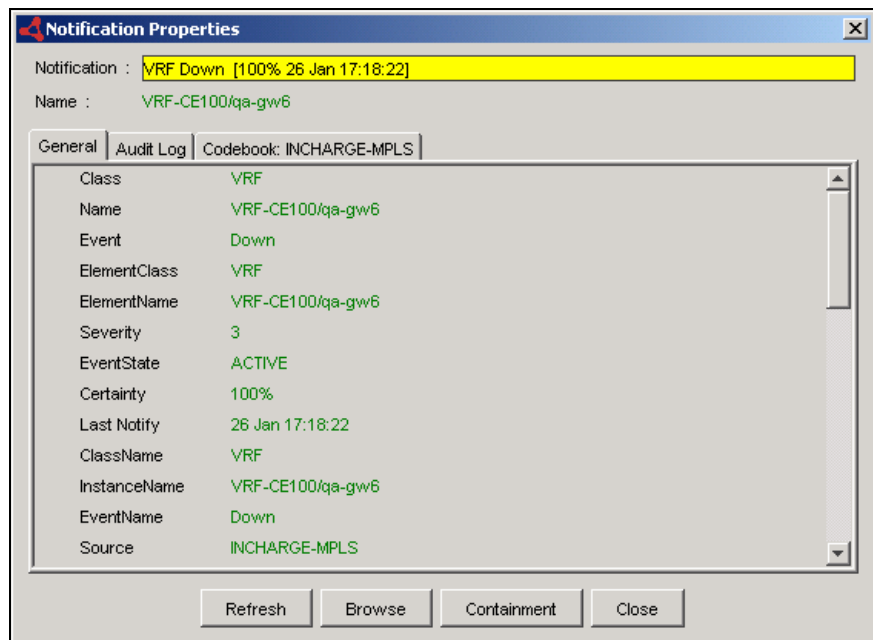


Figure 9: Notification Properties Dialog Box Showing a Down VRF

The VRF in Figure 9 either has no associated interfaces or has one or more associated interfaces and all of them are down.

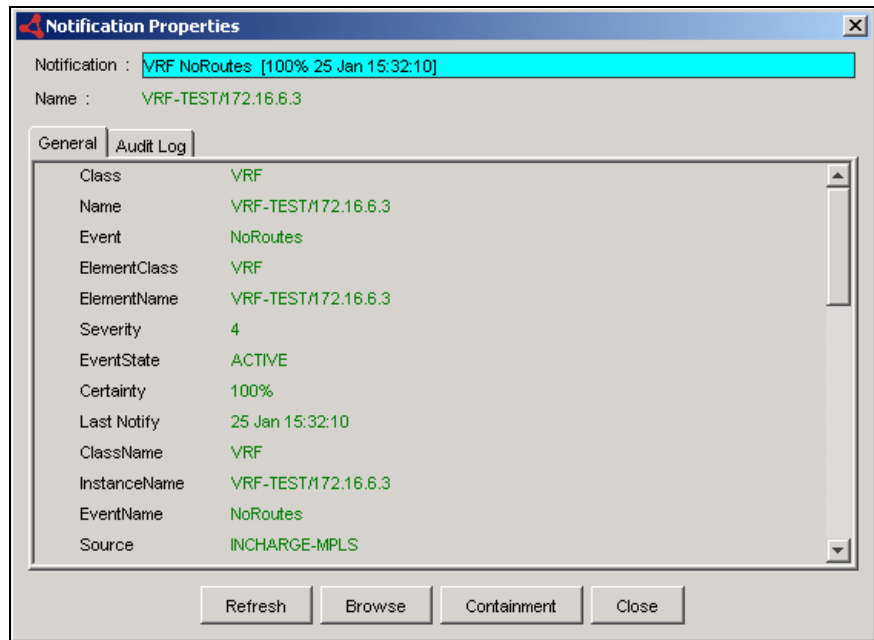


Figure 10: Notification Properties Dialog Box Showing a Routeless VRF

The VRF in Figure 10 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

MPLS Manager sends a streamlined copy of its MPLS topology—instances of the MPLSService, LSP, LSPHop, VPN, and VRF classes—to the Global Manager. The Global Manager combines this topology with the underlying transport network topology received from 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, or VRF) 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 to display a map for the element, or right-click an MPLS element and select an MPLS map type (LSP, LSP Hops, or VPN) 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, or VPN) from the pop-up menu.










MPLS Topology Map Graphical Representations

An MPLS topology map contains router, LSP, LSP hop, VPN, and VRF elements; and 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 14 identifies and describes the default icons 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 icons with other map icons that are preferred by your organization. In that case, use *Map > Map Legend* to see the definitions of your map icons.

Table 14: Default Nodes and Edges for MPLS Topology Maps

ICON / VISUAL INDICATOR	DESCRIPTION
	Icon—standard router icon with PE inscription—represents a Provider Edge (PE) router and the MPLS service element associated with the PE router.
	Icon—standard router icon with P inscription—represents a Provider (P) router and the MPLS service element associated with the P router.
	Icon—standard router icon with CE inscription—represents a Customer Edge (CE) router and the MPLS service element associated with the CE router.
	Icon represents a VPN.
	Icon represents a VRF.
	Solid line can represent a physical connection, a logical IP connection, a logical VLAN connection, a membership, or a group relationship.
	Jagged line can represent a network connection between routers or a virtual link between a VRF and a CE router.
	Solid black line with arrow can represent a dependency or an LSP in <i>No Highlight LSP</i> mode.
	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, 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 and VRFs. There are three types of maps that focus on MPLS elements:

- *LSP map*—available to all MPLS elements
- *LSP Hops map*—available to all MPLS elements; default map for LSPs
- *VPN map*—not available to LSPs; default map for VPNs and VRFs

In general, the display of an LSP, LSP Hops, or VPN map depends on the source element from which the map is launched.

LSP Map

The LSP map (Figure 11) shows the LSP connectivity between PE routers.

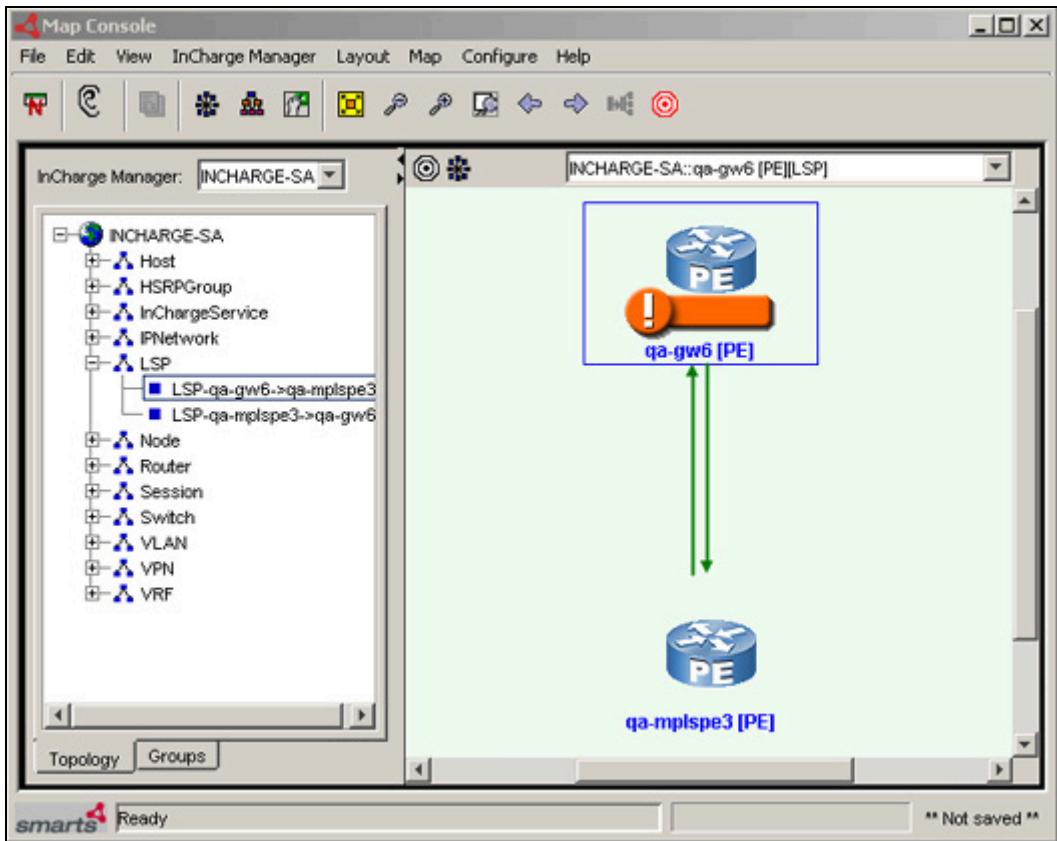


Figure 11: LSP Map—Example

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.

The display of an LSP map depends on the source element from which the map is launched:

- When launched from a PE router, the LSP map shows all the LSPs that either originate or terminate at the PE router.
- When launched from an LSP, the LSP map shows the LSP and the two PE routers associated with it.

- When launched from a VPN, the LSP map shows the PE routers and the LSPs belonging to the VPN in either the Full-Mesh or Hub-and-Spoke configuration.
- When launched from a VRF, the LSP map shows all the LSPs used by the VRF to communicate with its peer VRFs.

LSP Hops Map

The LSP Hops map (Figure 12) shows the intermediate LSP hops that comprise an LSP.

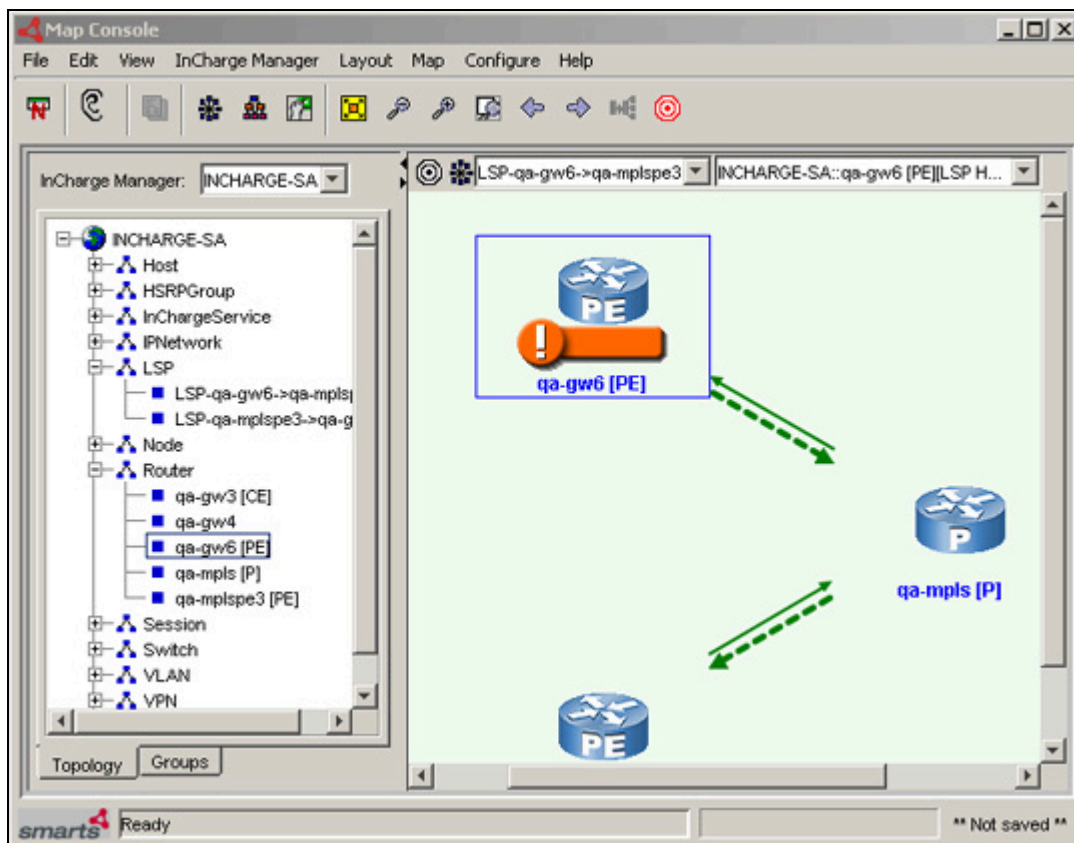


Figure 12: LSP Hops Map—Example

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 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

The VPN map (Figure 13) 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.

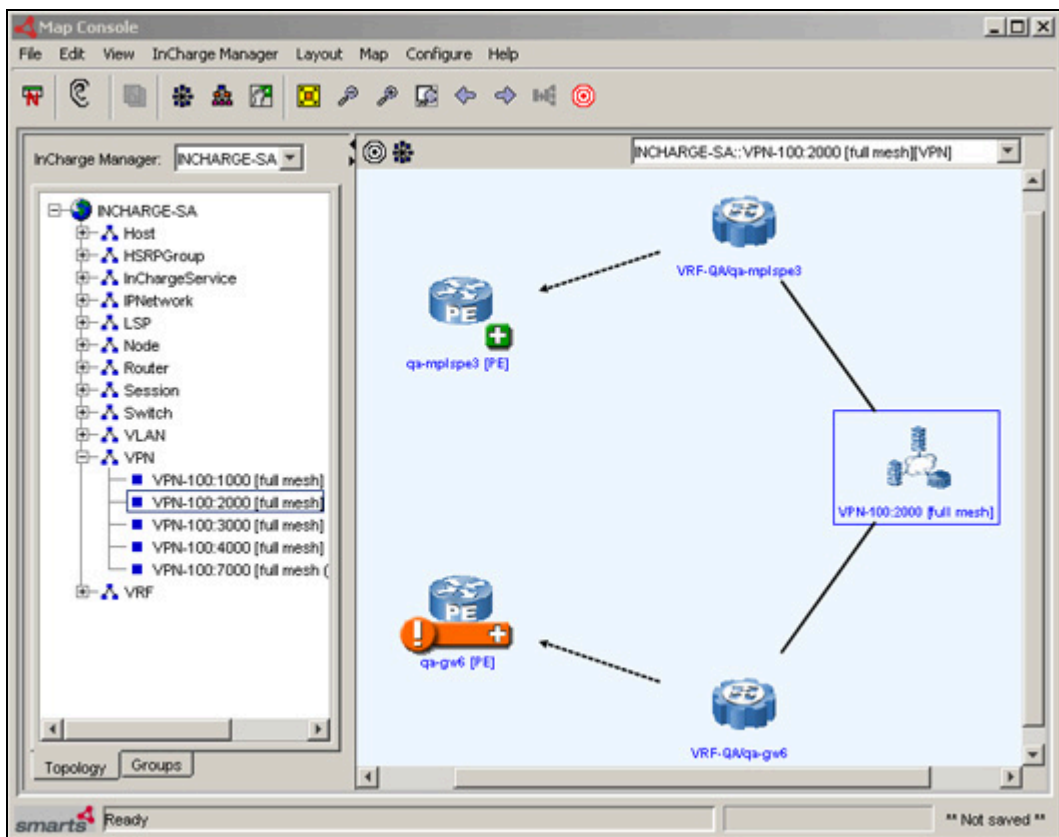


Figure 13: VPN Map—Example

The display of a VPN map depends on the source element from which the map is launched:

- When launched from a PE router, the VPN map shows all the VRFs hosted by the PE router.
- When launched from a CE router, the VPN map shows the VRF to which the CE router is virtually connected.
- When launched from a VPN, the VPN map shows all the VRFs that are part of the VPN and the PE routers hosting the VRFs.
- When launched from a VRF, the VPN map shows the VRF, the PE router hosting the VRF, and the CE routers to which the VRF virtually connects.

A map icon, such as a PE router icon, having a plus sign (+) next to it represents a container of additional connectivity information. To show the additional connectivity information, right-click the icon and choose **Expand Node** from the pop-up menu. For example, “expanding” a PE router container shows all the VRFs hosted by the PE router.

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 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 *InCharge MPLS Manager User’s Guide for Cisco ISC Adapter*. For information about the Business Impact Manager, see the *InCharge Service Assurance Manager User’s Guide for Business Impact Manager*.

5

Customizing MPLS Polling

MPLS Manager monitors the MPLS network by sending SNMP polls to the discovered VRF elements and comparing the polling results to threshold values that define acceptable and unacceptable levels of operation. MPLS Manager uses the thresholds values, in conjunction with events from Availability Manager, to monitor VRF availability and to detect VRF space overload.

Using the Polling and Thresholds Console, shown in the figure [Polling and Thresholds Console—Example](#) on page 43, you can customize the SNMP polling for MPLS Manager by modifying default polling groups or by creating new polling groups. The polling groups periodically collect the data needed by MPLS Manager to determine the availability of the VRF elements in the managed MPLS environment. The polled data serves as input to MPLS Manager correlation analysis.

For a description of SNMP polling for correlation analysis, see [Polling for Analysis](#) on page 63.

Polling Groups and Settings

The polling groups and settings for 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.

Polling Groups

Currently, MPLS Manager provides one default polling group named *VRFs*, which has a single default setting named *VRF SNMP Setting*. The *VRFs* polling group is used by MPLS Manager to monitor VRF instances.

The *VRFs* polling group accepts only elements of class *VRF*, so only *VRF* elements become members of the *VRFs* polling group. You can specify additional matching criteria (*VRF* attribute values) for the *VRFs* polling group to further limit its membership, as explained in [Editing Matching Criteria](#) on page 46.

Polling Settings

Currently, MPLS Manager provides the following polling settings:

- VRF SNMP Setting (default)
- VRF External Setting

These settings are mutually exclusive, meaning that one or the other (but not both) can be specified for an MPLS polling group.

VRF SNMP Setting

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

Table 15 lists the VRF SNMP Setting parameters.

Table 15: VRF SNMP Setting Parameters and Their Values

PARAMETER	VALUES	DESCRIPTION
AnalysisMode	ENABLED, DISABLED Default: ENABLED	Enables or disables MPLS Manager analysis; if analysis mode is disabled, no polling is performed.
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 allowed for the first poll request before it times out. Successive retries use longer times.

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

VRF External Setting

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

Table 16 lists the VRF External Setting parameters.

Table 16: 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.

Threshold Groups and Settings

Currently, there are no threshold groups and settings for MPLS Manager. 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 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, *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 *InCharge Service Assurance Manager Configuration Guide*.

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

- 1** Attach to the Domain Manager with the Global Console. 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. Likewise, when the Thresholds tab is selected, the console displays threshold groups. Threshold groups are not available to 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 14 provides an example of a Polling and Thresholds Console attached to an MPLS Manager instance named INCHARGE-MPLS.

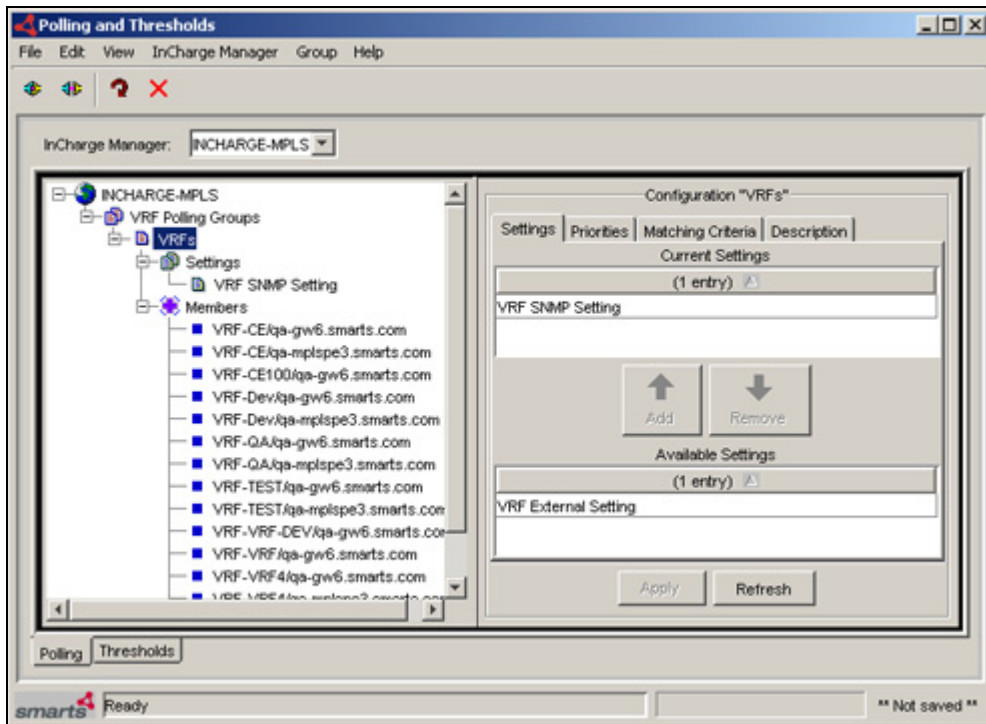






Figure 14: Polling and Thresholds Console—Example

Polling and Thresholds Console Toolbar Buttons

The toolbar of the Polling and Thresholds Console provides quick access to the commands described in Table 17.

Table 17: 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 Polling Groups and Settings

The configuration of a Domain Manager, such as MPLS Manager, applies polling parameters to defined sets of managed elements.

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

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

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

Modifying the Properties of a Group

A polling 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 are 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 settings determine what polling parameters are applied to the managed elements that are members of the group.

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

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

Priority and matching criteria 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 managed elements. 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 [Wildcards](#) on page 65) 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. When more than one matching criterion is specified, a managed element must match all criteria to become a member of the group.

Thus, for an MPLS polling group, which may only contain VRF elements, you can specify additional matching criteria (VRF attribute values) to restrict which managed VRF elements become members 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—for MPLS Manager, the group matches all managed VRF elements. Priority determines whether the group contains members.

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 Polling Groups

Creating a new polling group enables you to customize the polling 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 45, and for information regarding groups, see [Modifying the Properties of a Group](#) on page 44.

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.

A

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.

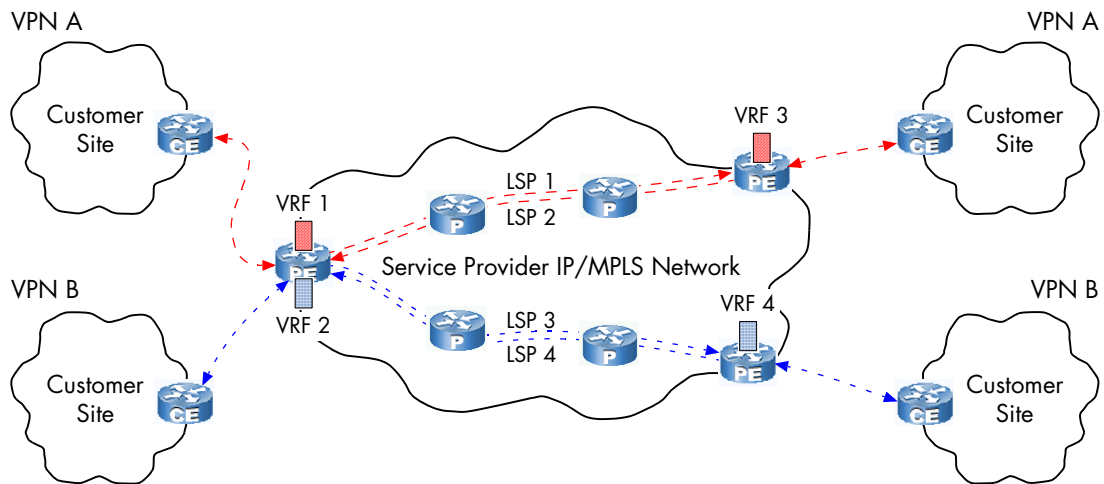


Figure 15: MPLS-Capable IP Network and MPLS VPNs

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) routers and PE routers, may be operated by regional Internet Service Providers (ISPs), local network operators, or even private companies.

- 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 an 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 Router

Customer Edge Router. A routing 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 2 or Layer 3 VPN.

- CR-LDP

Constraint-based Routing Label Distribution Protocol. A 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.

-
- **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.
 - **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 in this list.

- LDP
Label Distribution Protocol. A label signaling protocol used to advertise labels between PE and P routers to establish, maintain, and remove LSPs.
- 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 in this list.
- LSP
Label Switched Path. A concatenation of switch 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 in this list.
- LSP Segment
One hop between MPLS-enabled (PE/P) routers. An LSP tunnel 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 in this list.
- 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, 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.
 - **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 by applying MPLS labels to establish tunneled paths through non-IP networks such as Frame Relay, ATM, and Ethernet.
 - **MPLS FIB**

MPLS Forwarding Information Base. See MPLS Forwarding Table in this list.
 - **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, to labels, and used by P routers to rapidly switch data traffic through the MPLS network.
 - **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 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. (Currently, MPLS Manager does not support MPLS Layer 2 VPNs.)

- 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 RFC-2547bis VPN routing.

An ingress PE router examines the incoming packet's IP address, determines a route, assigns an LSP, and attaches a label (attaches two labels for a VPN packet). An egress PE router removes the label from the outgoing packet's IP address and forwards the packet to its destination via standard IP routing. The fact that a label temporarily exists 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 a 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.

- Route Distinguisher

An 8-byte value placed in front of a BGP IPv4 network route advertisement to identify the VRF to which the particular 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.

- 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.

- **RSVP-TE**

Resource Reservation Protocol with Traffic Engineering extensions. A 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.
- **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.
- **VPN Packet**

Virtual Private Network Packet. An RFC 2547bis packet transferred between the endpoints of the VPNs via MPLS. IP packets are labeled by the ingress PE router, based on the appropriate FEC, and forwarded to the proper P router. The only difference between an RFC 2547bis VPN packet and any other MPLS packet is that the VPN packet carries two labels: The outer label is used to route the packet through the MPLS network to the appropriate egress PE router, and the inner label is used to deliver the packet to the correct end user.
- **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 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*.

- 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.

- VRF

VPN Routing and Forwarding. A forwarding table used by the PE routers to establish 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

MIBs Polled

The MPLS Manager supports SNMP V1 and SNMP V2C for the MIB objects defined in the MPLS-LSR MIB, the MPLS-VPN MIB, and the jnxVpnMIB.

MIB Support for Cisco Devices and Juniper M/T Devices

For Cisco devices, MPLS Manager polls the following standard MIBs to gather discovery data for MPLS and VPN elements:

- MPLS-LSR MIB
- MPLS-VPN MIB

For Juniper M/T devices, MPLS Manager polls the following enterprise MIB to gather discovery data for MPLS and VPN elements:

- jnxVpnMIB

MPLS Manager reads certain MIB objects in these MIBs to discover MPLS and VPN elements and to instrument the elements to obtain their status. If MIB discovery is not successful, MPLS Manager attempts Command Line Interface (CLI) discovery to obtain the MPLS and VPN topology information, as described in [CLI Commands](#) on page 61.

MIB Support for Juniper ERX Devices

Because Juniper ERX devices and ERX virtual routers do not currently support MPLS or VPN MIBs, the MPLS Manager attempts CLI discovery at the very beginning of the discovery process to obtain the MPLS and VPN topology information, as described in [CLI Commands](#) on page 61.

C

CLI Commands

The MPLS Manager uses specific Command Line Interface (CLI) commands to obtain information regarding MPLS-enabled interfaces on devices, the MPLS forwarding tables, VPN Routing and Forwarding (VRF) instances configured on the devices, and Multiprotocol Border Gateway Protocol (MBGP) sessions between the devices and other MPLS-enabled devices. Assuming that CLI discovery is enabled, as described in the *InCharge MPLS Manager Configuration Guide*, MPLS Manager uses a CLI discovery probe during the discovery process to invoke the CLI commands.

CLI Support for Cisco Devices and Juniper M/T Devices

For Cisco and Juniper M/T devices, the MPLS Manager initially attempts to obtain the MPLS and VPN topology information using SNMP. If the requested MIBs are not available on the device (because the device operating-system version does not support the MIBs), the MPLS Manager uses the CLI discovery probe to establish Telnet sessions with the devices and to issue specific CLI *show* commands to obtain the MPLS and VPN topology information.

CLI Support for Juniper ERX Devices

For Juniper ERX devices and ERX virtual routers, which currently do not support MPLS or VPN MIBs, the MPLS Manager uses the CLI discovery probe at the very beginning of the discovery process to establish Telnet sessions with the devices and to issue specific CLI *show* commands to obtain the MPLS and VPN topology information.

D

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.

E

Wildcards

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 18. 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 19.

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 18: 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>serveró.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 * matching no characters. Similarly, it matches strings like 9x, with * matching the trailing x. However, it does not match 11, 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.

Table 18: Basic Wildcard Patterns (continued)

CHARACTER	DESCRIPTION
	Matches alternatives. For example, "ab bc cd" without spaces matches exactly the three following strings: "ab", "bc", and "cd". A as the first or last character of a pattern accepts an empty string as a match. Example with spaces "ab bc" matches the strings "ab " and " bc".
\	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 19: Compound Wildcard Patterns

CHARACTER	DESCRIPTION
&	"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. Example: *NY*&*Router* matches all strings which contain NY and also contain Router. 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. Example: *A* *B*&*C* matches strings that contain either an A or a B, and also contain a C.
~	"Except" for a compound wildcard pattern (opposite function of &). If a component basic wildcard pattern is preceded by ~, it <i>must not</i> match. Example: 10.20.30.*~10.20.30.50 matches all devices on network 10.20.30 except 10.20.30.50. 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.

Index

A

- Adapter
 - Cisco ISP 19
- Adding or removing a setting 45
- Adding or removing matching criteria 46
- Analysis 4
- Attributes
 - LSP 10
 - LSPHop 11
 - LSPInSegment 12
 - LSPOutSegment 13
 - MPLService 8
 - RouteTarget 19
 - VPN 15
 - VRF 16
- Availability Manager 1, 2

B

- BASEDIR x

C

- Changing matching criteria 46
- Changing priority of a group 45
- Changing setting parameters 47
- Cisco ISC adapter 19
- CLI commands 61
- Copying a group 48
- Creating a group 48

D

- DeviceType attribute 8
- Discovery 2
- Domain Manager Administration Console 41

E

- Edges
 - Table of 32
- Element
 - Assigning to groups 44

G

- Global Console 4, 29
 - Domain Manager Administration Console 41
 - Map Console view 29, 31
 - Notification Log Console view 29
 - Notification Properties dialog 30
 - Polling and Thresholds Console 42
- Global Manager 1, 4
- Group
 - Assigning members 44
 - Changing priority 45
 - Copying 48
 - Creating 48
 - Definition of 39
 - Properties 44

I

- ifIndex attribute 13
- Impact event
 - LSP 24
 - Notification 24
 - Summary table 24
 - VPN Impacted 24
 - VRF Impacted 24

K

- Key attribute 13, 19

L

- Label attribute 11, 12, 13
- Label Switched Path (LSP) 9
 - LSP 9
 - Attributes 10
 - LSPId 10
 - Name 10
 - Impacted 24
 - LSP Hops Map 36
 - LSP Map 34
 - LSPHop 10
 - Attributes 11
 - Label 11
 - LSPId 11

- Name 11
- LSPIId attribute 10, 11, 12, 13
- LSPIInSegment 12
 - Attributes 12
 - Label 12
 - LSPIId 12
 - Name 12
- LSPOutSegment 13
 - Attributes 13
 - ifIndex 13
 - Key 13
 - Label 13
 - LSPIId 13
 - Name 14
 - NextHopIP 14

M

- Map Console
 - Icons and indicators 32
 - Type of map 34
- Maps
 - LSP 34
 - LSP Hops 36
 - MPLS topology 31
 - Opening 32
 - Type of 34
 - VPN 37
- Matching
 - Pattern 65
- Matching criteria
 - Adding or removing 46
 - Changing 46
- MaxRoutes attribute 16
- MidRouteThreshold attribute 16
- Misconfiguration events
 - RouteTarget
 - Misconfiguration 7, 19
 - Summary table 7
 - VRF
 - Down 7, 17
 - MaxRoutesReached 7, 18
 - NoRoutes 7, 17
 - WarningThresholdCrossed 7, 18
- Monitoring 4
- MPLS Manager 1
- MPLSService 8
 - Attributes 8
 - DeviceType 8
 - Name 8

- NumberOfVRFs 8
- Supports_LSR_MIB 9
- Supports_VPN_MIB 9
- TotalVRFRoutes 9

N

- Name attribute 8, 10, 11, 12, 14, 15, 17, 19
- NextHopIP attribute 14
- Nodes
 - Table of 32
- Notification Log Console 4
- Notification Properties dialog 30
 - Opening 30
- Notifications
 - Impact events 24
 - Misconfiguration events 7
- NumberOfRoutes attribute 17
- NumberOfVRFs attribute 8

O

- Opening a Map 32
- Opening a Notification Properties dialog 30
- Operator
 - Wildcard 66

P

- Pattern 65
- Pattern matching 65
- Penultimate hop popping 10
- Polling
 - Groups 40
 - Settings 40
 - SNMP 4, 39, 64
- Polling and Thresholds Console 41
 - Layout 42
 - Polling tab 42
 - Thresholds tab 42
 - Toolbar buttons 43
- Polling tab 42
- PollingInterval
 - VRF SNMP setting 40
- Priority
 - Changing 45
- Provisioning system adapter 19

R

- Removing or adding a setting 45
- Removing or adding matching criteria 46

Restoring default values of a setting 47

Retries

 VRF SNMP setting 40

RouteDistinguisher attribute 17

RouteTarget 18

 Attributes 19

 Key 19

 Name 19

 Misconfiguration 7, 19

S

serverConnect.conf 42

Setting

 Adding or removing 45

 Changing parameters 47

 Definition of 39

 Restoring default values 47

SNMP

 Polling 4, 39, 64

Supports_LSR_MIB attribute 9

Supports_VPN_MIB attribute 9

T

Technical Support xii

Thresholds tab 42

Timeout

 VRF SNMP setting 40

Topology attribute 15

TotalVRFRoutes attribute 9

V

Virtual Private Network (VPN) 14

VPN 14

 Attributes 15

 Name 15

 Topology 15

 Impacted 24

VPN Map 37

VPN Routing and Forwarding (VRF) 15

VRF 15

 Attributes 16

 MaxRoutes 16

 MidRouteThreshold 16

 Name 17

 NumberOfRoutes 17

 RouteDistinguisher 17

 VRFName 17

 Down 7, 17

 Impacted 24

 MaxRoutesReached 7, 18

 NoRoutes 7, 17

 WarningThresholdCrossed 7, 18

VRF External setting

 MaxRoutes 41

 MidRoute Threshold 41

VRF polling

 VRF External setting 41

 VRF SNMP setting 40

VRF SNMP setting

 AnalysisMode 40

 PollingInterval 40

 Retries 40

 Timeout 40

VRFName attribute 17

W

Wildcard 65

 Chart of operators 66

