

**Oracle Communications IP Service Activator™  
Version 5.2.4**

# **Configuring VPN Services**

**Configuring MPLS and Layer 2 Martini VPNs,  
Transparent LAN Services and Circuit Cross Connects**

Fourth Edition  
December 2008

**ORACLE®**

Copyright © 1997, 2008, Oracle. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

# Contents

**Preface ..... ix**

- Before contacting Oracle Global Customer Support (GCS) ..... ix
- Contacting Oracle Global Customer Support (GCS) ..... x
- Downloading products and documentation ..... x
  - Downloading a media pack ..... x
- Service Activator publications ..... xi

**Chapter 1 Setting Up RFC2547 MPLS VPNs ..... 1**

- Introduction ..... 2
- Planning MPLS VPNs ..... 3
  - VPN topology ..... 3
  - Routing protocols ..... 5
  - Options for handling VRF tables ..... 6
  - Management VPNs ..... 12
  - Applying QoS or measurement to a VPN ..... 12
- Setting up an MPLS VPN ..... 13
- Before setting up MPLS VPNs ..... 13
  - Manual pre-configuration ..... 13
  - Setting up domain parameters ..... 15
  - Discovering the network and assigning roles ..... 19
  - Assigning devices to proxy agents ..... 20
  - Setting devices to Managed ..... 20
- Setting up customers ..... 20
- Setting up sites ..... 21
  - Associating a physical component with a site ..... 22
  - Setting up PE-CE routing parameters ..... 24
  - Setting up private and public addresses for PE interfaces ..... 31

---

Configuring IP unnumbered Private PE IP addresses .....	32
Setting advanced VRF table options .....	36
Setting network and aggregate statements .....	38
Specifying metrics for route redistribution .....	38
Setting up OSPF properties for a site .....	43
Setting up RIP properties for a site .....	45
Setting up EIGRP properties for a site .....	45
Setting up the VPN .....	46
Applying QoS to CE devices or SAA to a VPN .....	46
Creating an MPLS VPN .....	46
Linking sites to the VPN .....	49
Using an RD number per VPN or per site .....	49
Specifying a hub site .....	52
Creating a VPN map .....	53
Listing the sites in a VPN .....	54
Implementing the VPN .....	55
Viewing implemented VPNs .....	57
<b>Chapter 2 Setting Up Layer 2 Martini VPNs .....</b>	<b>59</b>
Layer 2 Martini VPNs .....	59
Benefits of Layer 2 Martini VPNs .....	59
Technical description of Layer 2 Martini VPNs .....	60
Layer 2 Martini VPN devices and data types .....	62
Overview of Layer 2 Martini VPN creation .....	65
Discovering devices and assigning roles for VPN setup .....	66
Creating a customer .....	67
Checking Interface Capabilities .....	68
Completing other pre-configuration for Layer 2 Martini VPNs .....	68
Provisioning endpoints (VC IDs) for a Martini L2 connection .....	70
Creating a Layer 2 Martini VPN .....	71
Modifying Layer 2 Martini VPNs .....	72
<b>Chapter 3 Setting Up Transparent LAN Services .....</b>	<b>75</b>
Overview .....	76

VC-LSPs .....	76
Transport LSPs .....	77
802.1 support .....	77
Mapping Frames to the TLS .....	77
Planning the TLS .....	79
Creating a TLS .....	80
Port-based TLS .....	80
Port and VLAN-based TLS .....	82
Creating a TLS to support a Stacked VLAN .....	88
Before setting up a TLS .....	89
Manual pre-configuration .....	89
Discovering the network and assigning roles .....	89
Assigning devices to proxy agents .....	90
Setting devices to Managed .....	90
Setting up customers .....	90
Setting up the TLS .....	90
Creating a TLS .....	91
Setting up layer 2 sites .....	94
Associating a physical component with a layer 2 site .....	96
Linking sites to a TLS .....	96
Applying rate limiting to a layer 2 site .....	97
Implementing the TLS .....	99
Viewing implemented TLSs .....	101
<b>Chapter 4 Setting Up Point-To-Point Connections .....</b>	<b>103</b>
Overview .....	104
Layer 2 switching CCC .....	105
MPLS tunneling CCC .....	106
Manual pre-configuration .....	106
Provision sub-interfaces for CCCs .....	107
Creating a CCC .....	108
Associating interfaces with a CCC .....	110
Implementing CCCs .....	110

---

Viewing implemented CCCs .....	112
Deleting CCCs .....	112
Applying QoS to CCCs .....	113
PHB groups .....	113
Access rules .....	114
<b>Chapter 5 IPsec VPNs .....</b>	<b>115</b>
First Time User Setup .....	115
Oracle Database Login .....	115
Service Activator Login .....	116
Creating Global Default Options .....	117
Creating Customer Default Options .....	120
Configuring IPsec and IP Tunnels .....	124
Service Activator GUI Setup .....	124
Multiple Interface Overview .....	126
Main Summary Screen Overview .....	128
Creating an IP Tunnel VPN .....	129
Creating an IPsec Tunnel VPN .....	133
Creating an IPsec over IP Tunnel VPN .....	139
Administration Tasks .....	140
Viewing an Existing IPsec/ IP Tunnel VPN .....	140
Deleting an IPsec/IP Tunnel VPN .....	141
Cleaning up Driver Scripts .....	141
VRF-Aware IPsec .....	142
Example VRF-Aware IPsec implementation .....	142
<b>Appendix A Setting Up Management and Customer VPNs .....</b>	<b>145</b>
Introduction .....	146
Steps to configure the VPNs .....	147
<b>Appendix B Example Usage: Interface Configuration Management Module .....</b>	<b>151</b>
Example: Creating Cisco Serial Sub-interface .....	152
Interface Creation Policy Registration .....	152

---

Creating a Sub-interface .....	153
Removing a Sub-interface .....	155
Example: Cisco Serial Sub-interface Decoration .....	156
Interface Creation Policy Registration .....	156
Decorate an existing sub-interface .....	157
Removing Sub-interface decoration .....	157
Example: Cisco Channelized Serial Interface Creation .....	158
Pre-requisites .....	158
Interface Creation Policy Registration .....	158
Configuring the E3 Controller .....	161
Create the Channellized Serial Interface .....	163
Removing channellized interface .....	165
Hints and Tips .....	166
Device Level Configuration Policy Workaround .....	166
Interface Creation and Decoration .....	166
Using Collection Based Configuration Policies .....	166
Centralized Management of Configuration Policy GUI Extensions .....	168
<b>Appendix C Oracle Communications Policy Services and IPSA</b>	
<b>Integration .....</b>	<b>171</b>
Integration Architecture .....	172
Detailed Behaviour .....	172
Graphical User Interface Presentation .....	173
<b>Index .....</b>	<b>175</b>





## Preface

The [Configuring VPN Services](#) guide is intended for operations managers, administrators, and activation technicians who are responsible for setting up VPN services using Oracle Communications IP Service Activator. It explains how to configure RFC2547 MPLS-based VPNs, Layer 2 Martini VPNs, Transparent LAN Services, Metro Ethernet VLANs, Circuit Cross Connects, and Internet Protocol Security (IPsec) tunnels.

For information on configuring MPLS Label Switched Paths (LSPs), refer to the MPLS LSP Module online help.

This guide consists of the following chapters:

- [Chapter 1: Setting Up RFC2547 MPLS VPNs](#) explains how to set up Multi-Protocol Label Switching (MPLS) VPNs within Cisco, Juniper or Unisphere based networks.
- [Chapter 2: Setting Up Layer 2 Martini VPNs](#) explains how to set up Layer 2 Martini tunnels on Cisco and Juniper M-series based networks.
- [Chapter 3: Setting Up Transparent LAN Services](#) explains how to configure a Transparent LAN Service.
- [Chapter 4: Setting Up Point-To-Point Connections](#) explains how to configure point-to-point connections, including configuring Circuit Cross Connects on Juniper devices.
- [Chapter 5: IPsec VPNs](#) explains how to configure point-to-point IPsec and IP/GRE tunnels.
- [Appendix A](#) explains how to set up management and customer VPNs correctly.

## Before contacting Oracle Global Customer Support (GCS)

If you have an issue or question, Oracle recommends reviewing the product documentation and articles on MetaLink in the Top Technical Documents section to see if you can find a solution. MetaLink is located at <http://metalink.oracle.com>.

In addition to MetaLink, product documentation can also be found on the product CDs and in the product set on Oracle E-Delivery.

Within the product documentation, the following publications may contain problem resolutions, work-arounds and troubleshooting information:

- Release Notes
- Oracle Installation and User's Guide
- README files

## Contacting Oracle Global Customer Support (GCS)

You can submit, update, and review service requests (SRs) of all severities on MetaLink, which is available 24 hours a day, 7 days a week. For technical issues of an urgent nature, you may call Oracle Global Customer Support (GCS) directly.

Oracle prefers that you use MetaLink to log your SR electronically, but if you need to contact GCS by telephone regarding a new SR, a support engineer will take down the information about your technical issue and then assign the SR to a technical engineer. A technical support representative for the Oracle and/or former MetaSolv products will then contact you.

Note that logging a new SR in a language other than English is only supported during your local country business hours. Outside of your local country business hours, technical issues are supported in English only. All SRs not logged in English outside of your local country business hours will be received the next business day. For broader access to skilled technical support, Oracle recommends logging new SRs in English.

Oracle GCS can be reached locally in each country. Refer to the Oracle website for the support contact information in your country. The Oracle support website is located at <http://www.oracle.com/support/contact.html>.

## Downloading products and documentation

To download the Oracle and/or former MetaSolv products and documentation, go to the Oracle E-Delivery site, located at <http://edelivery.oracle.com>.

You can purchase a hard copy of Oracle product documentation on the Oracle store site, located at <http://oraclestore.oracle.com>.

For a complete selection of Oracle documentation, go to the Oracle documentation site, located at <http://www.oracle.com/technology/documentation>.

## Downloading a media pack

### To download a media pack from Oracle E-Delivery

1. Go to <http://edelivery.oracle.com>.
2. Select the appropriate language and click **Continue**.

3. Enter the appropriate **Export Validation** information, accept the license agreements and click **Continue**.
4. For **Product Pack**, select **Oracle Communications Applications**.
5. For **Platform**, select the appropriate platform for your installation.
6. Click **Go**.
7. Select the appropriate media pack and click **Continue**.
8. Click **Download** for the items you wish to download.
9. Follow the installation documentation for each component you wish to install.

## Service Activator publications

The Service Activator documentation suite includes a full range of publications. Refer to the Service Activator *Release Notes* for more information.



## Chapter 1

# Setting Up RFC2547 MPLS VPNs

This section explains how to set up RFC2547bis MPLS VPNs (Virtual Private Networks) within Cisco, Alcatel 7670, Juniper M-series or Juniper E-series based networks.

It includes the following:

- Planning MPLS VPNs
- Important points that you need to consider before setting up a VPN, including pre-configuring routers and setting up relevant parameters
- Setting up customer sites and associating them with the appropriate routers/interfaces
- Creating VPNs and applying them to the network
- Viewing implemented VPNs

## Introduction

An IP virtual private network (VPN) is a means of creating a private network over a shared IP infrastructure. A VPN enables a secure, private connection between a number of geographically remote customer sites. VPNs can be used to implement corporate intranets, linking remote offices or mobile workers, and extranets, extending the services to customers, suppliers or other communities of interest.

Service Activator supports VPNs based on RFC2547bis, a widely supported IETF standard. This is an architecture based on using MPLS to forward packets over the backbone and using BGP (Border Gateway Protocol) to distribute routes.

MPLS VPNs are based on Layer 3 connectionless technology. The primary advantage of this is that MPLS VPNs are much more scalable than other IP VPN technologies, as there is no need to create a full mesh of tunnels or permanent VCs between all sites in the VPN. Deploying and managing an MPLS VPN is therefore more straightforward than the VC-based or IP tunneling options.

Security and privacy within an MPLS VPN is achieved by limiting the distribution of routing information to all members of the VPN. Routes to VPN sites are only advertised to members of the same VPN, and are not shared with core devices or devices outside the VPN.

Service Activator allows you to set up VPNs quickly and easily by defining appropriate customer sites and specifying how they are linked together into VPNs. The relevant routers throughout the network are then configured automatically. VPN membership can be updated by adding and deleting customer sites when required.

**Note:** Minimal manual configuration of routers is initially required.

Service Activator supports MPLS VPNs implemented within Cisco, Alcatel 7670, Juniper M-series, Juniper E-series, Huawei AR or Huawei NE networks. We strongly recommend that you read the relevant device driver or device cartridge guide before setting up an MPLS VPN. Each device driver guide describes how MPLS VPNs are implemented on a specific device type and provides example configurations.

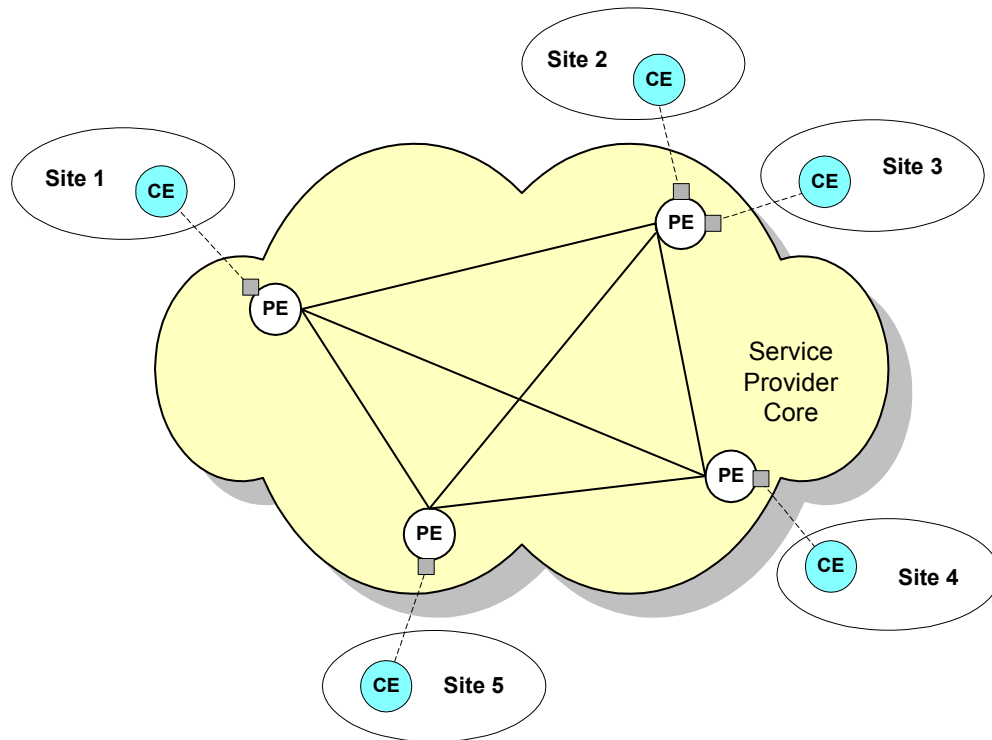
## Planning MPLS VPNs

There are a number of points you need to consider before setting up VPNs.

### VPN topology

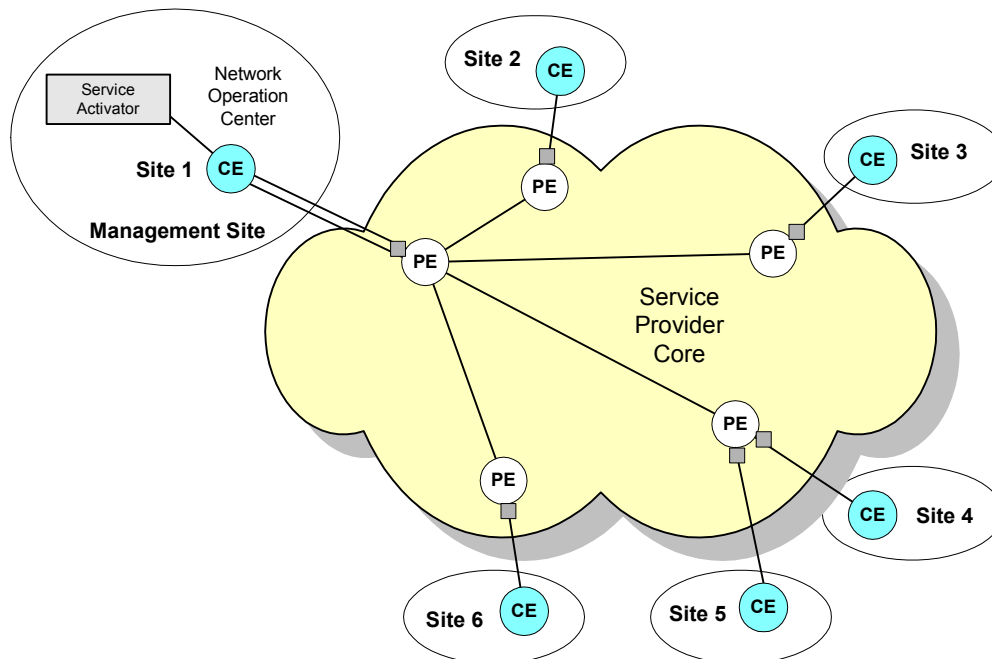
There are two types of VPN topology — fully-meshed and hub and spoke.

In a fully-meshed VPN, each customer site can communicate with all other sites. For example:



A site can be in any number of fully-meshed VPNs.

In a hub and spoke VPN, one or more sites act as a controlling or management interface. For example:



Sites 2, 3, 4, 5 and 6 are only aware of Site 1, while Site 1 can communicate with all other sites.


Each site must be defined as a hub or a spoke. A site may be a member of several hub and spoke VPNs and act as a hub in one VPN and a spoke in another. There may be more than one hub site in a hub and spoke VPN or, alternatively, all sites may be defined as spokes.

A management VPN is a special type of hub and spoke VPN that provides connectivity to CE devices and reduces the risk of connectivity loss. Create a management VPN if you intend to apply QoS or Service Assurance Agent (SAA) to a hub and spoke or fully-meshed VPN. For more information, see [Management VPNs on page 12](#).



## Routing protocols

The various routers within the VPN need to communicate using a routing information distribution protocol that defines who can talk to whom. The following diagram illustrates the routing protocols required in an MPLS VPN:



PE routers connected to other routers within the same VPN communicate using BGP, an IP routing protocol that defines how routes can be distributed. A BGP autonomous system is a collection of networks under a common routing strategy. Each autonomous system is identified by an autonomous system number (ASN), which is required when running BGP. Routers that belong to the same autonomous system and exchange BGP updates run internal BGP (iBGP). In order to configure BGP, you need to assign an ASN to each customer site network, in addition to the service provider core network. BGP is typically set up and configured automatically on PE devices by Service Activator.

Each PE router also needs to communicate with its external neighbors – the CE routers to which it is connected.

In addition to static routing, Service Activator supports the following protocols:

- External BGP (eBGP)
- RIP (Routing Information Protocol)
- OSPF (Open Shortest Path First)
- EIGRP
- EBGP & OSPF
- EBGP & RIP
- EBGP & EIGRP
- None

Service Activator supports only eBGP, RIP and static routing for PE-CE connectivity on Juniper E-series devices and Cisco devices.

Service Activator can configure static routes between the PE and CE router and these may be used in combination with any routing protocol. For multi-homed sites, you can configure a static route per PE interface. For eBGP and static routing, you need to set up appropriate routing parameters so that the PE router can be configured appropriately.

You can specify a metric to associate with routes as they are distributed from the PE-CE routing protocol into other IGPs and BGP and vice versa. For more information, see [Specifying metrics for route redistribution on page 38](#).

In cases where an MPLS VPN includes devices in more than one ASN, you should configure Service Activator not to configure iBGP peering sessions (the default). Service Activator will then adapt to use the ASNs already on the devices. If there is only a single ASN involved in an MPLS VPN configuration, then Service Activator uses the ASN specified in the User Interface.

## Options for handling VRF tables

A VPN must be secure and maintain data separation — it must prevent communication between sites that are not in the same VPN. One way to do this is to ensure that VPNs have their own routing tables in the PE router, so a customer site that belongs to a VPN can access only the set of routes contained in that routing table.

Each PE router maintains a number of separate forwarding tables known as VRF (VPN Routing and Forwarding instance) tables, and each site (i.e. each PE interface

or sub-interface connected to a CE device) must be mapped to one of those VRF tables.

Note that a VRF table does not necessarily correspond to a particular VPN. Its purpose is to hold the routes that are available to a particular site connected to a PE device. If a site is in multiple VPNs, the VRF table associated with that site contains routes within all the VPNs of which it is a member.

### Using VPN-wide or site-specific VRF details

By default, Service Activator automatically generates a site-specific VRF table name and Route Distinguisher (RD) for each site that participates in a VPN. In a typical MPLS VPN setup, the RD uniquely identifies the site.

However, you can override the Service Activator default by specifying at the VPN level that the same VRF table name and RD number is applied to all sites that participate in the VPN. You can choose whether to use Service Activator-generated values or specify your own VRF table name and/or RD number. Sites that participate in the VPN must be set to inherit VRF\RD details from the VPN.

A site may be set to inherit VRF / RD details and be a member of more than one VPN that specifies VPN-wide VRF / RD details. In this situation, Service Activator's default behavior is to generate VRF / RD details for the site to avoid any conflict. However, it is possible to specify that, where a site inherits VPN-wide VRF / RD details from multiple VPNs, user-defined details specified at site level are used instead.

Service Activator allows you the following options for defining the VRF table name and RD. You can:

- Specify that the same VRF table name and/or RD is used by all sites within a VPN.  
You can choose whether to accept the identifiers generated by Service Activator or create user-defined identifiers.
- Specify that each site has a manually-defined VRF table name and/or RD

Using a single RD number for all sites in a VPN is suitable only where a site belongs to one intranet VPN. If the site may become a member of an extranet VPN in the future, this method is not recommended.

- Specify that, if a site is member of only one VPN, its VRF table name and RD are derived from the VPN and, if a site is part of more than one VPN, its VRF table name and RD are manually defined at the site level.

For more information, see [page 36](#) and [page 49](#).

## VRF re-use or reduction

From a logical point of view, each VRF table maps onto an interface in a site (except with interface-less VRFs). However, having one VRF table for each PE interface in a site will create scalability problems on the router. Therefore, if multiple VRF tables contain exactly the same routing information (for example, if one site connects to two interfaces, or there are two sites that are members of the same VPN) and the routing protocol behaviors are identical or compatible between them, Service Activator normally reduces them to just one VRF table, in order to minimize resource usage. This is known as VRF reduction. The method Service Activator uses to perform this process depends on whether you are using system-generated or manually defined VRF table names.

Where system-defined VRF table names are used, the first VRF table creation request with a particular

When a new VRF table creation request is sent to the device driver, the driver first attempts to reduce the VRF table into an existing Service Activator provisioned VRF. If this is not possible, a new VRF is provisioned. In case of multiple concurrent VRF table creation requests, the driver evaluates the request in the order of increasing site RD numbers (for system-defined VRF table names) or in the ASCII sort order of the VRF table names (for user-defined VRF table names).

If the new VRF has the Force Install attribute set, then no reduction is performed

If the VRF which is a candidate for reduction with the new VRF was created with Shareable attribute turned off, no reduction is performed.

If VPN attributes or routing protocol attributes do not match or are incompatible no reduction is performed.

When a VRF is initially created, the site which triggered the creation of the new VRF is known as the 'VRF creator'. If other VRF tables are reduced into this VRF, a dependency on the VRF creator remains.

If the VRF creator site is taken out of the VPN, the VRF is de-provisioned. All remaining interfaces are reduced into a newly created VRF. One of these sites will be the new VRF creator site for this VRF in accordance with the ordering logic explained above.

If the VRF creator site is modified such that its VRF is no longer compatible with the other VRF tables which are reduced into it, the interfaces of the other sites are removed from the VRF before the VRF is modified to reflect the new VRF creator site properties. The removed interfaces are reduced into a newly created VRF. Again, of these interfaces' sites will be the new VRF creator site for this new VRF in accordance with the ordering logic explained above.

During the re-provisioning occurring as a result of either of the above scenarios, the interfaces moved to the new VRF temporarily lose connectivity and trigger routing re-convergence.

## Interface-less VRFs and Sites

Service Activator supports the indirect creation of interface-less VRFs and therefore interface-less Sites. An interface-less VPN site models a VRF on a router where no interface points to the VRF.

For complete details, refer to the topic **Interface-less VRFs and Sites** in the Service Activator Online Help.

### Service Application Points

When an interface-less Site and VRF are created, an object called a Service Application Point is modelled in the background and linked to the Site. The Service Application Point object behaves similarly to an interface (and has a role of Access for purposes of supporting the interface-less Site and VRF) but it is not accessible or modifiable through the GUI. The PE device is displayed in the Access Points folder for the site in order to represent the Service Access Point.

Note that Service Application Point objects are exposed in the EOM and are accessible through the OSS Integration Manager interface. Refer to the *OSS Integration Manager Guide* for details.

### Creating an interface-less Site and VRF

To create an interface-less VRF, first create a VPN Site. Then, drag the entire PE device into the site as you normally would an interface. This creates the Service

Application Point and causes the PE symbol to be displayed in the Access Points folder of the Site to represent it. (See below).

Once created, the Site can participate in VPNs in a similar fashion to site with an attached interface.

### Site Properties - VRF Router-ID attribute

The **router-id** attribute is available on the **Site** dialog box, **Site VRF** property page. It is applicable only for use with Alcatel 7670.

The router-id can be set on a VRF regardless of the routing protocol chosen for the site. However if OSPF or EBGp is chosen as routing protocol, a router-id **must** be set or an error is returned. This error indicates that the SiteAddress requires a router-id.

When a VPN is committed, the device driver attempts to *reduce* the number of VRFs. During this process Service Activator ignores all the router-id properties except for one — the selection of the router-id as discussed below. Once the VRF has been reduced, it is created on the device. Before setting the router-id parameter in the VRF, the driver creates a loopback interface and binds it to the VRF. The appropriate IP address will also be set on this newly created loopback interface.

If a VRF is initially created for one interface, this interface/SiteAddress will remain designated as *'the one containing the router-id'* for the entire life of that VRF regardless of the number of interfaces that are subsequently added to it. The router-id can be added, if desired, to the first interface bound (and committed) to a VRF. Using **Force Install** also ensures that an Interface/SiteAddress is the first one in its VRF.

Unless loopback interfaces are filtered out during the discovery process, subsequent device re-discoveries may find the newly created loopback interfaces and display them in the GUI. Do not attempt to use them.

Some operations such as setting a specific VRF name, route distinguisher, or DHCP Helper to a SiteAddress may force a VRF 'split'. In this case, the router-id requirements may no longer be met for the target configuration. An error is returned indicating which SiteAddress(es) require an additional router-id.

### Working with manually pre-configured VRF tables

If an MPLS VPN has already been manually configured on a network, Service Activator can work with the pre-configured VRF tables that exist on devices. You can choose how Service Activator handles these tables:

- Ignore – Service Activator leaves the pre-configured VRF table 'as is' and does not update it
- Assume control of the VRF table and preserve existing content – Service Activator controls and updates the VRF table but leaves pre-existing content
- Assume control of the VRF table but remove existing content – Service Activator controls and updates the VRF table and removes any pre-existing content

Service Activator’s handling of pre-configured VRF tables is controlled by advanced VRF table options which can be defined per site or per VPN.

Site or VPN property	Manually pre-configured VRF		
	No control	Control & preserve content	Control & remove content
VRF table name	Use Service Activator VRF name	Specify pre-defined name	Specify pre-defined name
Route distinguisher	Use Service Activator RD*	Specify pre-defined RD	Use Service Activator RD†

\* If a manually pre-configured VRF table has an RD that matches the RD of another manually pre-configured VRF table that is subsequently controlled by Service Activator, the first pre-configured VRF table is replaced by the second pre-configured VRF table.

† There is a small possibility that Service Activator may generate the same RD as that of the pre-configured VRF table. In this case, the VRF table will be controlled by Service Activator and its contents preserved.

If you unmanage a device, then make a change to it (for example, a VRF parameter) and then commit these two actions in the same transaction, the change is saved in the Service Activator database, but is not propagated to the device.

If you remanage the device in a subsequent transaction, there is loss of synchronization between the device and the Service Activator database. If this occurs, repeat the change in a new transaction.

Remember, do not unmanage a device and make a change to the device in the same transaction.

Loss of synchronization also occurs if you make manual changes to a device which has been unmanaged, and is later remanaged. If this occurs, manually remove the incorrect configuration from the device.

## Management VPNs

If you need to have visibility of the customer's CE devices – for example, to apply QoS or measurement at the CE devices – you should set up a management VPN. The management VPN provides connectivity to CE or shadow devices and avoids potential loss of connectivity.

In Service Activator, a VPN is always associated with a customer. We suggest you create a customer named 'Management' and create the management VPN beneath this. Then create the 'real' customer and create a hub and spoke or fully-meshed VPN.

For a summary of the steps involved in setting up a management VPN, see [Setting Up Management and Customer VPNs on page 145](#).

## Applying QoS or measurement to a VPN

Service Activator's policy management features allow you to apply QoS throughout the VPN. Network traffic can be divided into separate classes of service and allocated quality of service characteristics according to service level agreements with the customer. For more information see the *Network and SLA Monitoring Guide*.



When applying QoS to a VPN, note that policy that has been applied to a customer is inherited to a site, unless the site is a member of a VPN to which a policy element of the same type has been applied. In this case, the policy applied to the VPN overrides that applied to the customer.

Measurement can also be applied to the VPN using SAA. SAA is a Cisco technology that enables you to apply measurement to point-to-point connections within an MPLS VPN or a measurement-only VPN. In a measurement-only VPN, devices are grouped solely for the purpose of applying measurement. For more information see the *Configuring SLA Monitoring*.

## Setting up an MPLS VPN

The following steps are involved in setting up a VPN within Service Activator:

- Set up domain-specific parameters ([Setting up domain parameters on page 15](#))
- Discover the network and assign appropriate roles ([Discovering the network and assigning roles on page 19](#))
- Set up the customer sites that are to be connected by a VPN, including information about their routing parameters ([Setting up customers on page 20](#) and [Setting up sites on page 21](#))
- Associate each site with the appropriate access interface on a PE (gateway) router ([Setting up PE-CE routing parameters on page 24](#))
- Set up a VPN object and link the appropriate customer sites to it by dragging and dropping ([Creating an MPLS VPN on page 46](#))

Minimal manual configuration of routers is initially required. This is detailed below.

## Before setting up MPLS VPNs

This section explains a number of steps you need to ensure are performed before setting up a VPN.

### Manual pre-configuration

Before setting up VPNs, some manual configuration of routers is required. For detailed information, see the appropriate device driver guide.

## PE routers and P routers

On all PE (gateway) and P (core) routers in the core network, you need to ensure that MPLS is enabled on core interfaces. An Interior Gateway Protocol (IGP) must be implemented in order to distribute IP routes.

Ensure also that all IP addresses are correctly assigned, and that a loopback interface is set up.

On Cisco devices, ensure that CEF (Cisco Express Forwarding) or dCEF (Distributed CEF) is configured, as it is a prerequisite for MPLS.

On Juniper M-series or Juniper E-series devices, you must configure LSPs between the loopback addresses of all PE and P devices.

### User-defined VRF tables, export maps and prefix filters

You can manually pre-configure PE routers with your own data to allow greater flexibility or special operational requirements to be implemented for your MPLS VPNs.

VRF tables can be manually configured to invoke particular VPN behavior and other special requirements.

This option is not currently available for Juniper M-series devices.

Export maps can be manually configured with the prefixes of VRF table routes which you want exported to other PE routers. A VRF table route is only exported to other PE routers if its prefix matches one of those specified in the export map. The exported routes are tagged with a route target specified by the export map. This ensures that routes are only advertised to sites that need to receive them.

You can also manually configure route maps, or policy statements on Juniper M-series routers, to filter the routes that are redistributed between the protocols used for PE-CE connectivity within the VPN.

Prefix filters can be manually configured with the prefixes of eBGP routes which you can specify are to be either accepted or denied by a VRF table. You can associate a prefix filter with incoming and/or outgoing routes.

To allow Service Activator to use manually pre-configured data, you must ensure that the default action for detecting manual configuration of network devices is set to either **Disabled** or **Warn** in the **Domain** dialog box.

## CE routers

The CE (access) routers at customer sites are not configured for VPNs by Service Activator, since they may not be under the control of the network service provider. Therefore they need to be manually configured. You need to ensure that BGP, RIP, OSPF or static routing is configured in order to advertise reachability information between the CE and the PE.

We also recommend that a loopback interface is set up on each CE router and configured to carry IP traffic.

## Setting up domain parameters

Before setting up the VPN, you should ensure the appropriate parameters are set up correctly on the tabs in the **Domain** dialog box (these may have been set up when the domain was created). These parameters define:

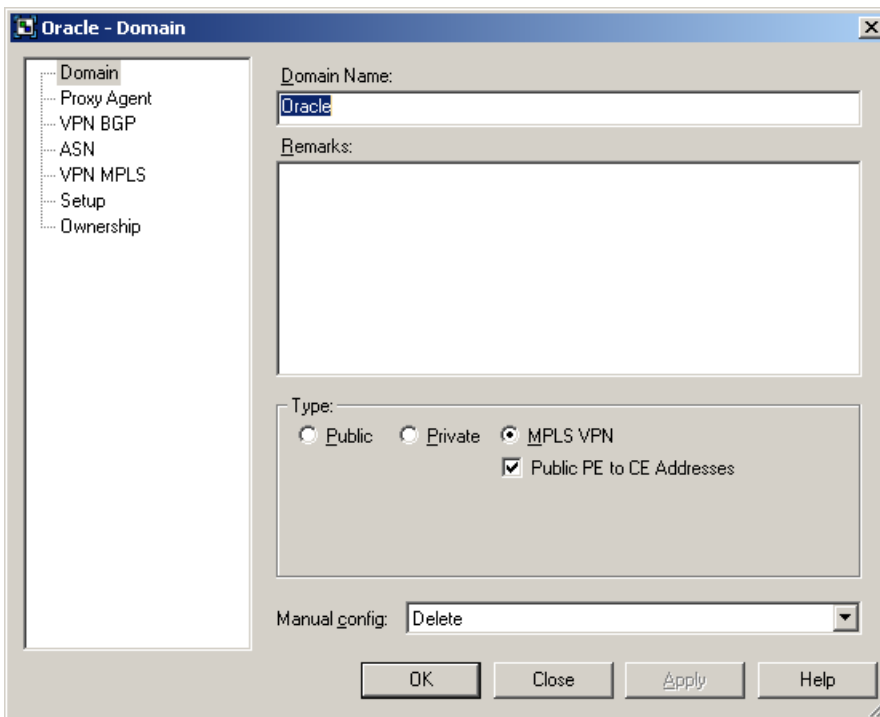
- Whether iBGP peering is configured by Service Activator
- How data in the AS\_PATH attribute is handled
- The number of alternative routes held by each PE to peer PE devices
- Whether BGP routes use the standard or extended community attribute or both
- Whether secure TCP connections are configured between iBGP peers using MD5
- Which interface is used as the loopback interface

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To check the domain parameters

1. Display the global setup window.
2. Select the **Domain** tab, right-click the relevant domain and select **Properties** from the pop-up menu.

The **Domain** dialog box for the selected domain is displayed.

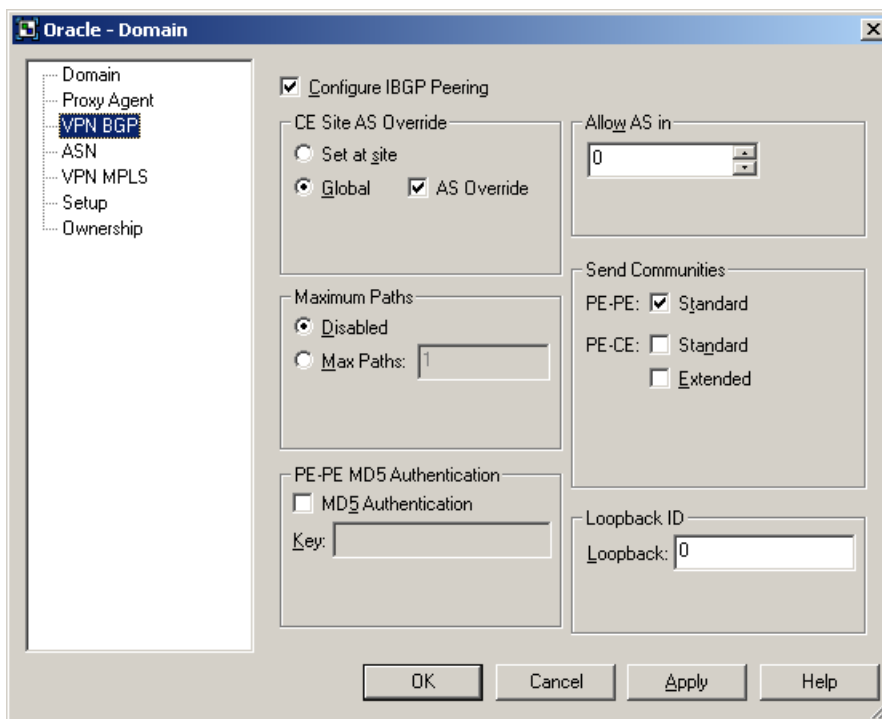


3. On the **Domain** page, check the following parameters:
  - The domain **Type** must be **MPLS VPN**.
  - The **Public PE to CE Addresses** checkbox specifies whether public or private addresses are used on PE routers. If selected, it indicates that the PE interface connected to a CE router uses a public address. If cleared, the PE interface is assumed to use a private address.

The **Manual config** field enables you specify how Service Activator handles manual pre-configuration it finds on the device. Note that Service Activator never deletes manually pre-configured VRF tables even if you select the **Delete** or **Warn and delete** options. If you select **Fail and don't delete**, no configuration is installed.

**Note:** The value set in the **Manual config** field on the **Domain** and **Device** property pages, including **Warn and Delete**, is ignored for **Juniper M-series devices**. The Juniper M-series Device Driver cannot be set up to monitor for or warn when changes to device configuration are made by other users. However, Service Activator can co-exist with manually applied configuration.

4. Select the **VPN BGP** page.



5. Specify values including **Configure iBGP Peering, CE Site AS Override, Allow AS in, Maximum Paths, Send Communities, PE-PE MD5 Authentication, and Loopback ID.**

**The default is for Service Activator not to configure iBGP peering.** If you leave this setting off, iBGP peering must already be configured correctly on your devices.

Note that if the **Configure iBGP Peering** checkbox is cleared, Service Activator leaves all iBGP configuration untouched. This means that whatever is installed will remain on the device, whether or not it was configured by Service Activator.

Setting	Description
Note that on Juniper M-series devices, the loopback ID must always be specified as 0 through the user interface.	

### Specifying a VRF route limit

You can specify the maximum number of routes that can be added to the VRF table maintained by each PE peer within the domain. You can also specify a threshold value at which a warning message is generated in the router log either when the limit is reached or when a percentage of the limit is reached. By default, no VRF route limit is specified.

You can also override the limit defined at the domain and specify a different setting for a particular site. For more information, see [Setting advanced VRF table options on page 36](#).

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

#### To specify a VRF route limit

1. Display the **Domain** dialog box and select the **VPN MPLS** page.
2. Select **VRF Route Limit** checkbox.
3. In the **Max Route** field, enter the maximum number of routes that can be added to a PE router's VRF table.
4. Select either **Warn only** or **Warn at n% from** the **Notification** combo box.

### Setting up the provider core ASN

To enable BGP communication throughout the VPN you must set the ASN for the domain.

#### To set the ASN for the domain

1. Display the **Domain** dialog box and select the **ASN** page.
2. Specify a value in the **Internal BGP ASN** field.

## Discovering the network and assigning roles

At this point you can run the discovery process to find all the P and PE routers in the network and include their details in Service Activator's database.

Run a discovery, as described in *Network Discovery and Basic Setup*.

Note the following:

- In an MPLS domain, the core provider network is assumed to use public addresses, and a hop count can be specified to discover further connected devices. All CE routers are assumed to use private addresses and an IP address or DNS name must be specified in order to discover them.
- All devices within the network must be correctly assigned system-defined roles, that is, PE routers must be classified as gateway devices, P routers classified as core devices and CE routers, if visible, as access devices. The recommended way of assigning roles is by means of role assignment rules, which automatically assign roles during device discovery. If you do not use role assignment rules you need to assign a role manually for each device. For more information, see *Network Discovery and Basic Setup*.

The system-defined Shadow role may be assigned to shadow routers associated with PE devices for the purpose of generating SAA measurement data. For more information, see the *Network and SLA Monitoring Guide*.

- All interfaces within the network must be correctly assigned system-defined roles:
  - On CE (access) devices, the interface connected to the PE device must be classified as an access interface. Interfaces connected to local segments must be classified as local interfaces.
  - On PE (gateway) devices, the interface connected to the CE device must be classified as an access interface. Interfaces connected to other PE devices or P (core) devices must be classified as core interfaces.
  - All interfaces on P (core) devices should be classified as core interfaces.

As for devices, we recommend you assign roles using role assignment rules. If you don't use role assignment rules you need to assign a role manually for each interface. For more information, see *Network Discovery and Basic Setup*.

If the PE interfaces connected to CE devices have private addresses, that is, if the **Public PE to CE Addresses** checkbox on the **Domain** dialog box is cleared, connectivity cannot be determined. Therefore, CE devices and access interfaces on PE devices will not automatically be connected to segments and interfaces will not automatically be assigned roles. The connections between the PE and CE can be applied manually by dragging one interface on to another on the map or set manually from the CE.

## Assigning devices to proxy agents

All devices that are to be managed by Service Activator must be assigned to a proxy agent. This is generally performed automatically during device discovery, but if devices are not assigned to the correct proxy agents you must assign them manually. For information on assigning devices to proxy agents, see *Network Discovery and Basic Setup*.

## Setting devices to Managed

All devices to be configured by Service Activator need to have their status set to Managed. When devices are first discovered, their status is Unmanaged.

### To set all devices to Managed

- Select **Manage All Devices** from the network or network map's pop-up menu.

### To set an individual device to Managed

- Select **Manage** from the device's pop-up menu.

The device's color changes to reflect its new status. A managed device is represented by a green icon, an unmanaged device is represented by a blue icon.

## Setting up customers

You must create a customer before you can create the sites and VPNs that feature in a service.

### To set up a customer

1. On the **Service** tab, right-click on the **Customers** folder and select **Add Customer** from the pop-up menu.

The **Customer** dialog box opens.

2. Enter values including **Customer name**, **Remarks**, and **Reference**.



3. Click **OK**.
4. Commit the transaction.

## Setting up sites

You need to set up a site for each member of a VPN. You create sites on a per-customer basis – you cannot create a site that is customer-independent. You can associate a site with more than one customer.

### To set up a site

1. You set up a site in the **Site** dialog box. You can open the **Site** dialog box on the **Service** tab:
  - On the **Service** tab, right-click on the relevant customer folder and select **Add VPN Site** from the pop-up menu.

The site object appears on the **Service** tab and the **Site** dialog box opens.

The screenshot shows the 'Site' configuration window. On the left is a tree view with the following structure:

- Site
- Addressing
- Connectivity
- VRF
  - VRF Advanced
  - VRF Maps
- Static Routing
- BGP Networks
- BGP Aggregate Address
- EBGP
  - Advanced
  - Route Maps
  - Dampening
  - Redistribution
- EIGRP
  - Redistribution
- OSPF
  - Summary Addressing
  - Redistribution
- RIP
  - Redistribution
  - Ownership

The main configuration area contains the following fields:

- Name: CustomerSite1
- Remarks: Customer Site1
- Account Ref: IPSAAC1
- Contact: Mr. Smith
- Address: Ottawa, Canada
- Tel: 565-1212
- Fax: 565-1212
- E-mail: techsupport@customer.com

Buttons at the bottom: OK, Cancel, Apply, Help.

2. On the **Site** page, specify an identifying **Name** for the site, and any additional comments. You can set up account and contact information if required, but this is optional.
3. Select **OK**.

Dragging a PE interface into a site adds the IP address of that interface to the **Addressing** page and avoids the need to enter the address manually. See [Setting up private and public addresses for PE interfaces on page 31](#).

## Associating a physical component with a site

Each site within a VPN must be defined by a physical network component with the exception of interface-less sites. (See [Interface-less VRFs and Sites on page 9](#) for more.) Typically a specific interface or device must be attached to the site in order to enable full routing information to be distributed throughout the VPN.

## Linking a PE access interface


You need to link the access interface of the appropriate gateway (PE) router to the site.

Once you link an interface with a site and define its routing details, Service Activator maintains that information even if you subsequently remove that interface from the site. This means that you can re-use routing details if you change the interface that is associated with a site.

For a multi-homed site, you can link several PE interfaces to the site.

### To link a PE access interface

1. Drag and drop the appropriate access interface on the gateway device on to the site.
2. If there are existing unused interface address details associated with the site, the **Site Addresses** dialog box opens:
  - **Create a new Site Address:** Create new public/private IP address details for the interface being linked.
  - **Attach to an existing Site Address:** Use existing address details – select an address from the list.



## Linking a CE router

If Service Activator has visibility of site CE routers, you should also link the CE device to the site directly. This is only required if the service provider is offering a fully-managed VPN service and has complete visibility of the customer's devices.

**To link the CE device**

- Drag and drop the access device on to the site.

**Virtual Tunnel Interfaces**

Juniper VT (virtual tunnel) sub-interfaces have VPN-MPLS capabilities which allows them to be added to VRFs.

**Setting up PE-CE routing parameters**

Support for particular protocols is device dependent. Consult the relevant device driver guide for details of the protocols supported.

You need to set up details of the routing protocol used between the PE router and the CE router. This can be eBGP, RIP or OSPF and different parameters are required for each. Static routing is also available and can be used in combination with other routing protocols. Static routes can be configured per PE interface for multi-homed sites.

For information on specifying the metrics to apply to routes as they are redistributed between protocols, see [Specifying metrics for route redistribution on page 38](#).

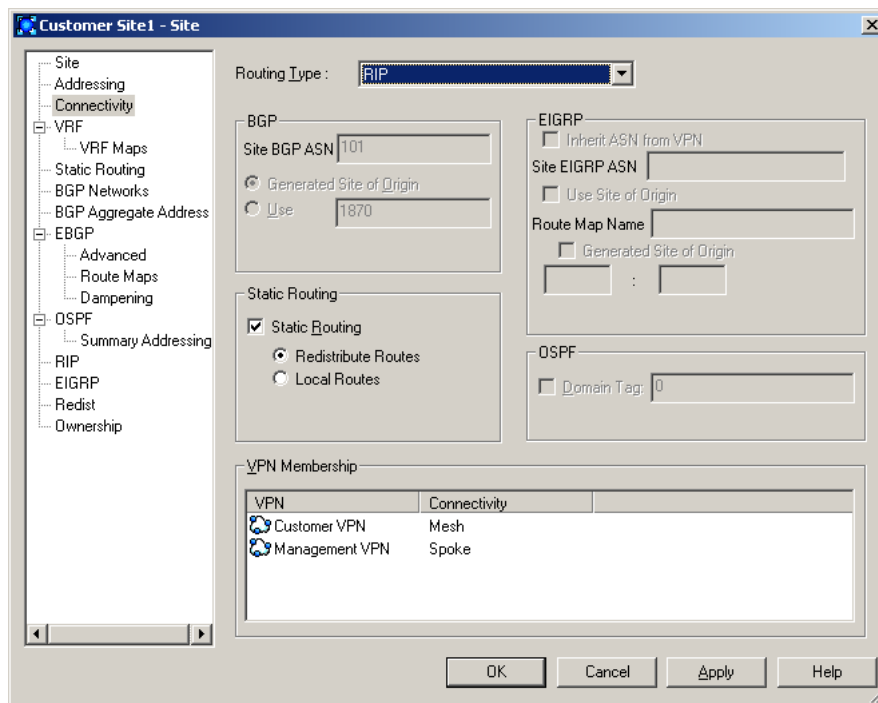
You can also select a routing protocol type of 'None' to distribute routes to VPN peers without configuring PE to CE routing.

Site of Origin (SOO) is configured automatically for sites that have more than one CE to PE connection where eBGP is used for PE-CE connectivity. For configuration details, see the appropriate device driver guide.

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

**To select a PE-CE routing parameter**

1. Display the **Site** dialog box and select the **Connectivity** page.



2. Specify the **Routing Type**. Select from **EBGP, RIP, OSPF, EBGP & OSPF, EBGP & RIP**, or **None**.
  - If **EBGP** is selected, you must supply additional routing parameters on the **EBGP** page and, optionally, the **EBGP Advanced, Route Maps**, and **EBGP Dampening** pages. See [Configuring eBGP parameters on page 26](#).
  - If **RIP** is selected, additional parameters are available on the **RIP** property page. See [Setting up RIP properties for a site on page 45](#).
  - If **OSPF** is selected, you may want to supply additional routing parameters. See [Setting up OSPF properties for a site on page 43](#).
  - If **EBGP & OSPF** is selected, both protocols are configured for the site. Static routing is still supported as well. Configure additional options for both protocols including route redistribution. See [Configuring eBGP parameters on page 26](#) and [Setting up OSPF properties for a site on page 43](#).
  - If **EBGP & RIP** is selected, both protocols are configured for the site. Static routing is still supported as well. Configure additional options for both protocols including route redistribution. See [Configuring eBGP parameters on page 26](#) and [Setting up RIP properties for a site on page 45](#).

- If **None** is selected, directly-connected VPN routes can be automatically redistributed to VPN members without path configuration by a routing protocol. This option should be used with option **Redistribute connected**.
  - Also select **None** if you wish to configure static routes between the PE and CE device instead of a routing protocol.
3. Specify values including **Static Routing, Redistribute Routes, Local Routes, Domain Tag, Generated Site of Origin** and **VPN Membership**.

For information on defining static routes, see [Configuring static routing parameters on page 30](#).

You can also specify metrics for route redistribution between protocols. For information, see [Specifying metrics for route redistribution on page 38](#).

### Configuring eBGP parameters

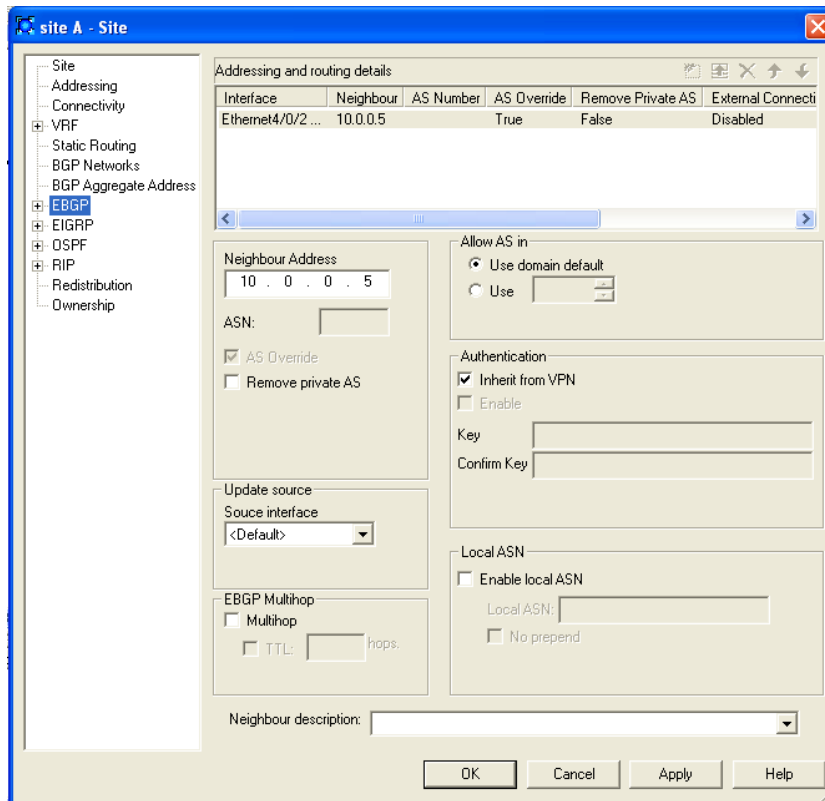
Service Activator provides control over eBGP configuration, enabling you to:

- Specify the number of times the same ASN can occur in the AS\_PATH attribute of a route prefix
- Specify the local preference for each interface in a multi-homed site
- Configure secure TCP connections between eBGP peers using authentication
- Define route dampening parameters
- Specify the number of alternative routes to the CE device that are maintained in the PE device's routing table

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

#### To specify eBGP parameters

1. In the **Site** dialog box, select the **EBGP** page.



2. Specify values including **Neighbour Address**, **AS Override**, **Remove private AS**, **Allow AS in**, **Authentication**, and **EBGP Multihop**.
3. In the **Update Source** drop-down, select default to use the PE Interface ID in the update-source interface, or enter your own description. You may leave this field empty to avoid propagating the update source interface.
4. Under **Local ASN**, select the check box **Enable local ASN** to configure Local Autonomous Systems Numbers. In the **Local ASN** field specify a numeric value for the Autonomous Systems path attribute. Select **No Prepend** if required.
5. In the **Neighbour Description** field, enter a description of the neighbor or select default. Leave this field empty to avoid generating the neighbor description command.

For more information on **AS Override**, see [page 17](#). For more information on **Allow AS in**, see [page 17](#).

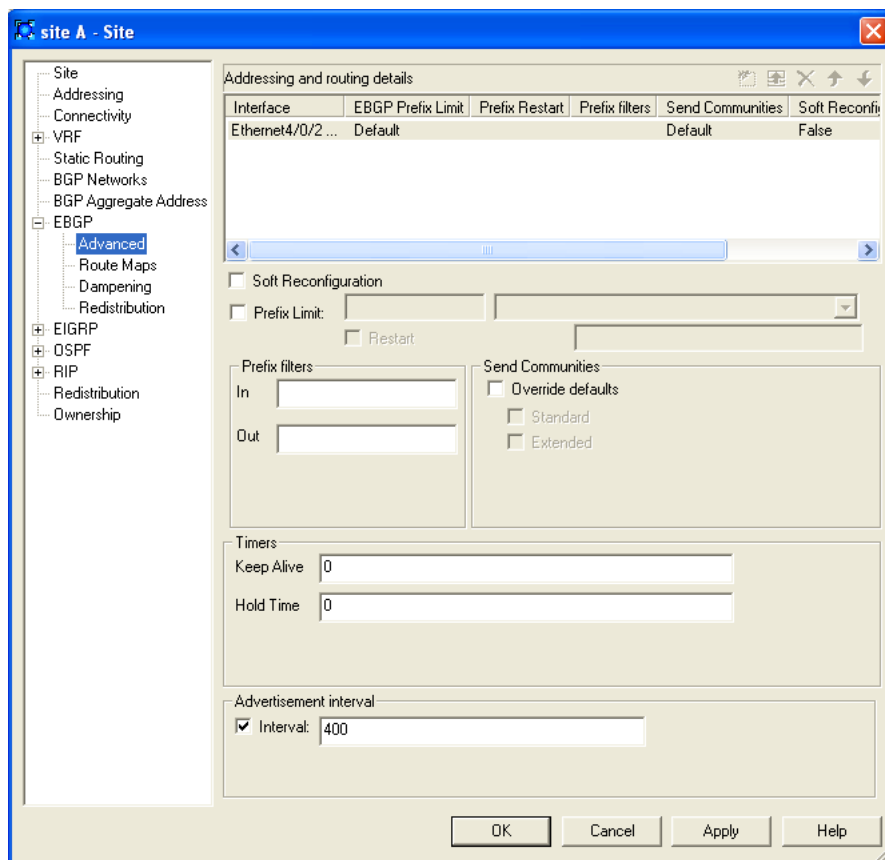
6. If multiple PE interfaces or sub-interfaces are associated with a site, specify settings for each listed interface or sub-interface.

Use the following button to edit an entry:



Set button: applies a change to the selected entry

- If you wish to define advanced eBGP parameters, select the **EBGP Advanced** page.



- Specify values for each listed PE interface or sub-interface including **Soft Reconfiguration, Prefix limit, Restart, Delay (minutes), Prefix filter, Send Communities, Timers, and Advertisement Interval.**

- Specify eBGP Route-map parameters, on the **EBGP Route Maps** page.

Service Activator supports inbound and outbound external BGP route-maps applied on a per-neighbor basis for the site.



**Note:** Use a naming scheme different from Service Activator's for external inbound and outbound route-maps. Service Activator will remove route-maps with the same names as those which it generates when **Use Autogenerated Route-map** in the **EBGP Route Maps** property page of the **Site** dialog box is enabled. This can also occur when the device is unmanaged and remanaged depending on the setting of the **Unmanage Action** attribute.

Route-map names specified on this property page are not validated against the names of route-maps provisioned on the router. You must correctly specify names of the externally defined route-maps.

10. Set the options for each PE interface listed in the **CE Addressing and routing details** listbox including **Inbound Route Map: Use External Route-Map, Use Generated Route-Map, Local Preference** and **Outbound Route Map: Use External Route-Map**.

11. If you wish to specify eBGP dampening parameters, select the **EBGP Dampening** page.

Route dampening is a mechanism that attempts to minimize instability by suppressing the advertisement of unstable routes. Penalties are applied when a route is withdrawn, re-advertised or changed. When a predefined penalty limit is reached, further advertisement of the route is suppressed. The penalty is reduced according to a defined 'half-life' setting, and once the penalty decreases below a limit, the route can be re-advertised.

EBGP dampening is supported on Alcatel 7670, Juniper M-series and Juniper E-series devices.

12. Select the **EBGP Dampening** checkbox and set options including **Decay Half-life, Reuse Threshold, Suppression Threshold, and Max Suppression Time**.

### Configuring a multi-AS site VPN

Service Activator supports VPNs which bridge more than one Autonomous System. However, in order to create this type of VPN, the eBGP and iBGP peering sessions must be configured manually.

When the VPN is configured on a device which already contains an ASN, and Service Activator is told not to configure iBGP peering, the iBGP configuration already on the device is left unaltered, as is the ASN.

In order to create a site which bridges multiple Autonomous Systems:

- Manually assign ASNs to the PE devices

- Manually configure iBGP peering between PEs in the same AS
- Manually configure eBGP peering between PEs in different AS clouds
- Set up Service Activator not to configure iBGP peering
- Perform the remainder of the VPN configuration in Service Activator as normal

### Configuring static routing parameters

You can specify how Service Activator configures the next-hop parameter in a static route. Choices include:

- IP Address and Interface
- IP Address Only
- Interface Only
- Null Interface

By default, Service Activator configures static routes with the interface name, next hop IP address and metric. Other configuration choices include:

- whether the next-hop-address is an address that is in the routing table and not the VRF table (**Global** checkbox)
- whether the static route will not be removed even if the interface shuts down (**Permanent** checkbox)
- whether a tag is to be associated with a static route, allowing it to be used by route map match statements controlling redistribution of routes (**Use this tag** field).

If you have manually configured static routes on the device, these routes are not removed by Service Activator provided the VRF table the routes are associated with is not controlled by Service Activator.

### Controlling redistribution of static routes

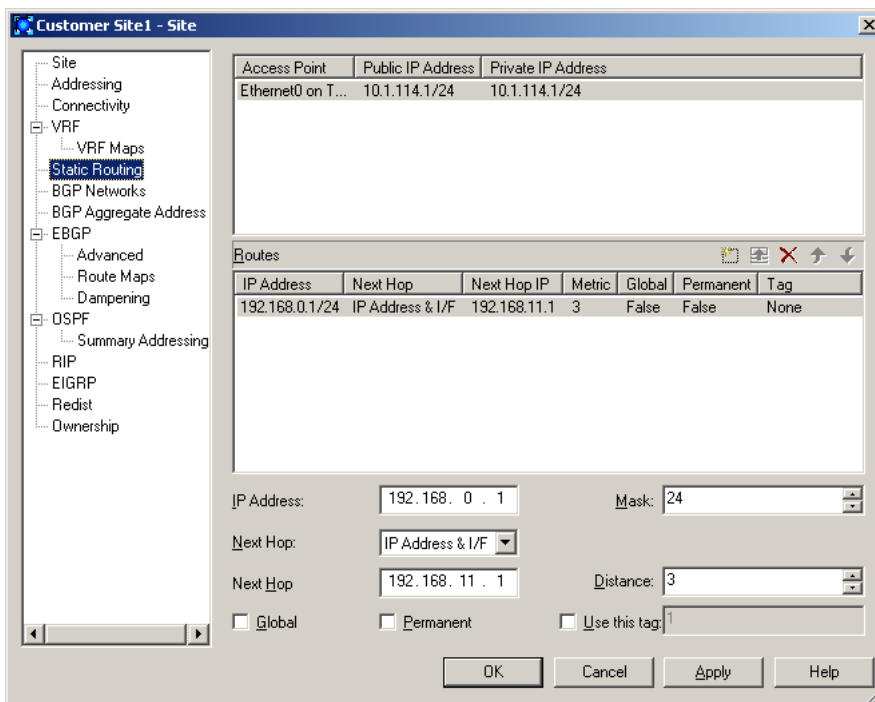
You can control whether or not static routes are redistributed into dynamic routing protocols. On the **Connectivity** property page of the **Site** dialog box, select **Redistribute Routes** to redistribute static routes. Select **Local Routes** to have static routes remain local.

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To specify static routing parameters

1. In the **Site** dialog box, select the **Static Routing** page.

**Note:** The fields on the Static Routing property page are disabled until the **Static Routing** checkbox on the **Connectivity** property page is selected.



2. Select a listed PE interface or sub-interface and specify values including **IP Address, Mask, Next Hop (Type), Next Hop (IP Address), Distance, Global, Permanent** and **Use this tag**.

## Setting up private and public addresses for PE interfaces

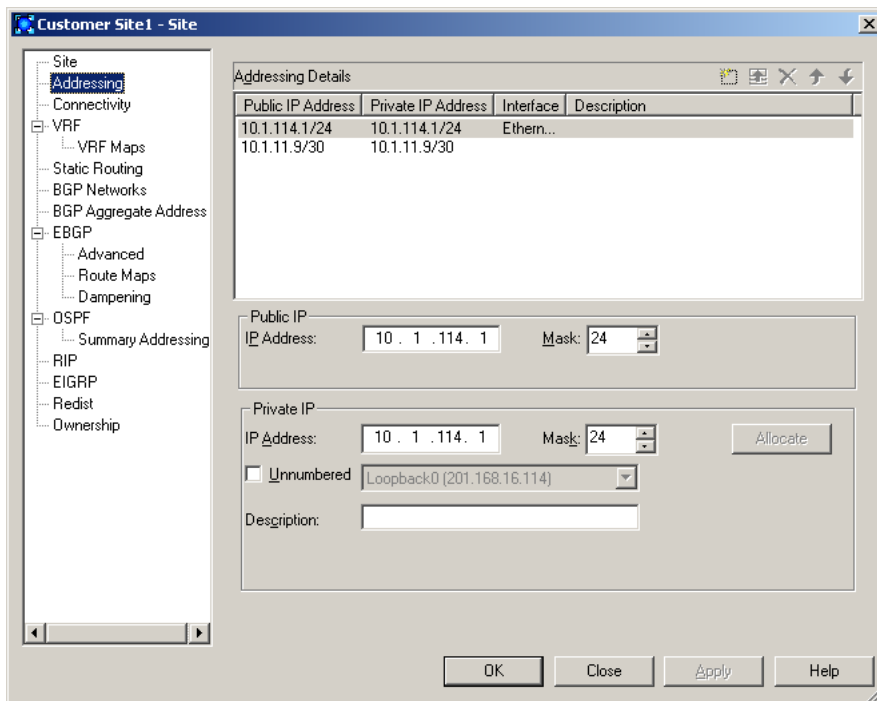
When an interface is added to a VPN it leaves the public IP space and becomes part of a private IP space. Therefore for the PE access interface you need to set up public and private addresses. The public address applies when the interface is outside the VPN, and the private address applies when it is within the VPN.

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To set up private and public addressing

1. Display the **Site** dialog box and select the **Addressing** page.

Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.



2. Select the desired interface in the **Addressing Details** list, and supply values including **Public IP and Mask, Private IP and Mask, Unnumbered** and **Description**.
3. If the INA Proxy server is used for assigning IP addresses, click **Allocate**. (The button is disabled if the INA Proxy server is not installed or not available.) This allows you to update the eBGP neighbor address for the site (as displayed on the **EBGP** property page).

See [Configuring IP unnumbered Private PE IP addresses on page 32](#) for details on support for IP Unnumbered interfaces. For more information on Policy Services and IPSA Integration, refer to [Appendix C on page 171](#).

Additional important notes on the **Description** field are available. For complete dialog box and property page descriptions, refer to the *Online Help*.

## Configuring IP unnumbered Private PE IP addresses

Service Activator supports IP unnumbered Private PE addressing for certain serial point-to-point IP interfaces in VPN sites on Cisco devices. This allows you to enable

IP on an interface and use it in a VPN without having to assign an explicit Private PE IP address and mask. Instead, a the IP address of loopback address from the device is used.

### **To configure an interface for IP unnumbered addressing**

IP unnumbered is configured on the Site properties - Addressing page property page by selecting the **Unnumbered** checkbox for the Private PE IP address for the interface in the Site. Then, select the loopback interface which will provide an outgoing IP address for the interface from the adjacent dropdown list. Alternatively, you can enter other interface details as well.

See [Setting up private and public addresses for PE interfaces on page 31](#) for details.

### **About routing protocols, VPNs and IP unnumbered**

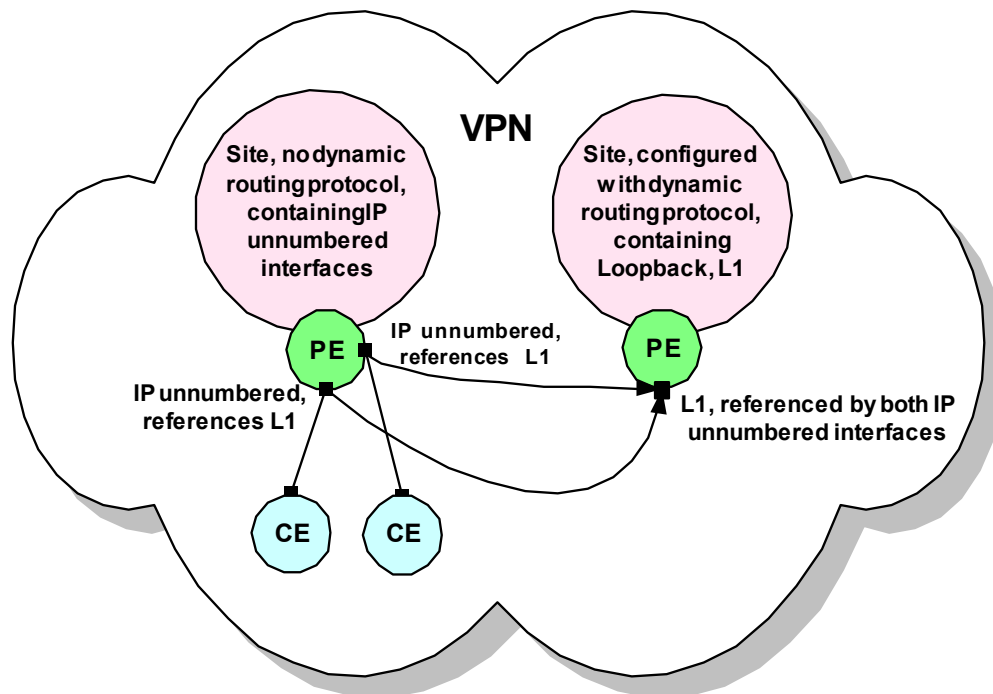
Note that dynamic routing protocol OSPF can not be used in a site if any of the PE interfaces linked to the site use IP unnumbered. (The Routing Protocol on the Site properties - Connectivity page must be set to **None**, **RIP** or **EBGP**) Static routing is permissible.

### **To indirectly associate a dynamic routing protocol to a site with IP unnumbered interfaces**

To configure indirect association of a dynamic routing protocol with a site containing IP unnumbered interfaces, link the loopback interface which the IP unnumbered interface references to another site which has the desired routing protocol configured. The dynamic routing protocol applied to the loopback in the other site will effectively apply to the IP unnumbered interface.

You can indirectly associate a dynamic routing protocol (OSPF, BGP and RIP for Cisco) with a site containing an IP unnumbered interface.

To do this, link a loopback interface to a second site which has the desired routing protocol configured on it. Then refer to that loopback from the IP unnumbered interface in the first site. See the illustration below.



### Error messages and warnings

An attempt to configure an interface's Private PE IP to IP unnumbered when the site has a routing protocol configured will trigger the following error: [572], A Routing Protocol can not be set if an interface uses IP Unnumbered.

Configuring IP unnumbered on an unsupported interface type triggers the following error: [587], IP Unnumbered is not allowed for this type of Interface.

If the loopback used for the IP unnumbered address is removed from a device, the following fault is raised in the fault pane: Site private PE ip unnumbered reference is invalid.

### Supported interface types for IP unnumbered

IP unnumbered can be configured on certain types of serial point-to-point interface (as reported during discovery):

- dsl (18)
- e1 (19)
- propPointToPointSerial(22)
- ppp (23)

- frameRelay (32)
- atm (37)
- sonnet (39)
- frameRelayService (44)
- v35 (45)
- ds0Bundle (82)
- async (84)
- atmSubInterface (134)
- rfc1483 (159)
- aal2 (187)

Attempting to configure IP unnumbered on an unsupported interface type will trigger an error stating IP Unnumbered is not allowed for this type of Interface.

**Restrictions, hints and tips**

Before a loopback interface can be assigned to an IP unnumbered interface configuration:

- it must be must be created
- an IP address assigned to it
- it must be discovered

Note that switching an interface from an explicit private PE IP address and mask to IP unnumbered and back does not affect its connectivity into the VPN - the interface continues to belong to the VRF.

You must have bridging configured for ATM devices to use IP unnumbered. The specific bridging protocols are device specific. For example:

Device	ATM Bridging Protocol
Cisco 827	RFC1483 Bridging and Integrated Routing and Bridging (IRB)
Cisco 6400 Asynchronous Transfer Mode	Routed bridge encapsulation (RBE)

**Sample configuration:**

```
interface ATM0/0/0.4 point-to-point
```

```
ip unnumbered Loopback1
no ip directed-broadcast
no ip route-cache
ATM route-bridged ip
PVC 4/100
encapsulation aal5snap
```

## Setting advanced VRF table options

Advanced options give you finer control of the VRF table and route handling within the VPN.

The following options are available:

- Specify whether the VRF table name and RD number are unique to the site or inherited from the VPN in which the site participates  
For information on using site-specific or VPN-wide VRF details, see [page 7](#).
- If using site-specific VRF and RD settings, specify whether to use the Service Activator default or a user-defined VRF table name.

By specifying a user or system-defined RD number and/or VRF table name, you can control how Service Activator handles manually pre-configured VRF tables. For more information, see [page 10](#).

For information on how user-defined VRF table names affect Service Activator's VRF reduction process, see [page 8](#).

- Control over the VRF table – specify whether every interface associated with a site has its own VRF table, or can be merged with the VRF table for another interface where routes in the table are identical.
- Specify a manually pre-configured route map that filters the routes exported from one site to PE peers within the VPN.
- Specify the maximum number of routes from a CE router that can be added to the VRF table – the maximum may be set for the site or inherited from the default defined for the domain. A warning message may be logged when the number of routes stored reaches a user-defined percentage of the maximum.

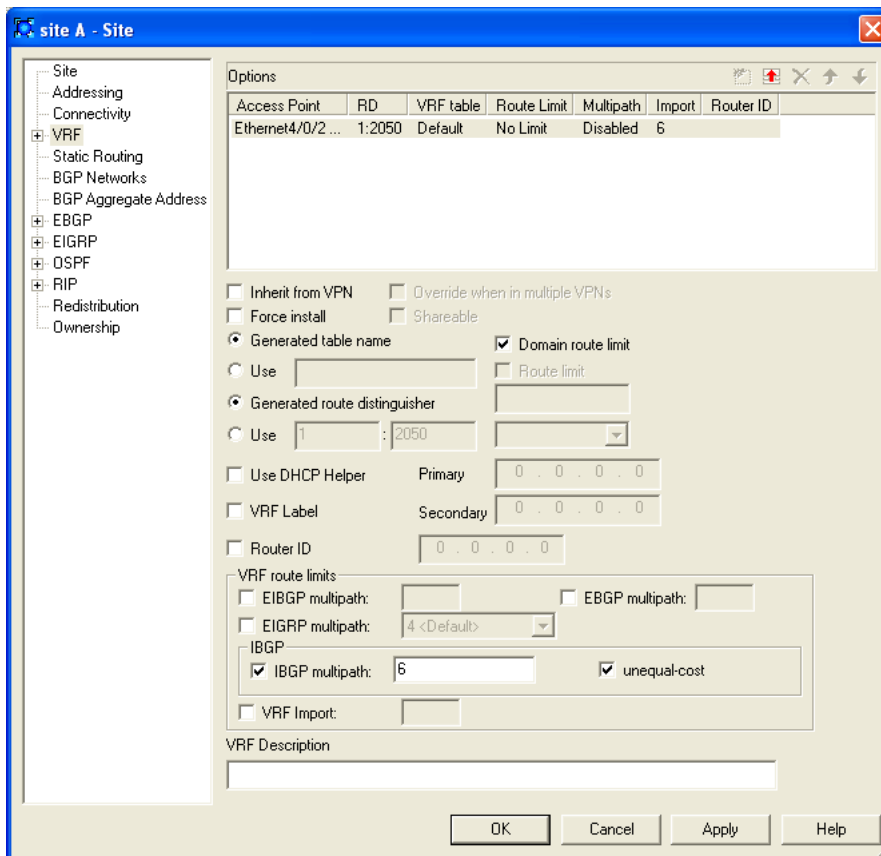
**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.



**To set advanced VRF table options**

1. Display the **Site** dialog box and select the **VRF** page.

Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.



2. Specify values for each listed PE interface or sub-interface including **Inherit from VPN**, **Override when in multiple VPNs**, **Force install**, **Shareable**, **Generated table name**, **Use**, **Generated route distinguisher**, **Domain route limit**, **Route limit**, **Use DHCP Helper**, **VRF Label**, **Router ID**, **EIBGP multipath**, **IBGP multipath**, **VRF Import**, and **VRF Description**.
3. Enable the check box **unequal-cost** to allow unequal cost load balancing by selecting iBGP paths that do not have an equal cost.

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To specify VRF import and export maps

1. Display the **Site** dialog box and select the **VRF Export** page.  
Any PE interfaces that you have linked to the site are listed. If no PE interfaces have been linked to the site, no addresses appear.
2. For each listed PE interface, specify values including Export map name and Import map name.

See [Setting Route Target numbers on page 47](#).

## Setting network and aggregate statements

In the **Site** properties dialog box, you can access the **BGP Networks** and **BGP Aggregate Address** property pages to set up network and aggregate statements. For complete dialog box and property page descriptions, refer to the *Online Help*.

Network statements are used to advertise networks to other routers. For the information to be advertised by BGP, a route to the specified network must be present in the routing table. This routing information can come from connected routers and dynamic routing or static routing sources.

Aggregate statements summarize routes into a single advertisement that is sent to BGP peers.

For more detailed, conceptual information on network and aggregate statements, see the Service Activator *Cisco Device Driver Guide*.

## Specifying metrics for route redistribution

You can specify the metric to apply to routes distributed from the selected PE-CE routing protocol into other Internal Gateway Protocols (IGPs) and BGP, and vice versa.

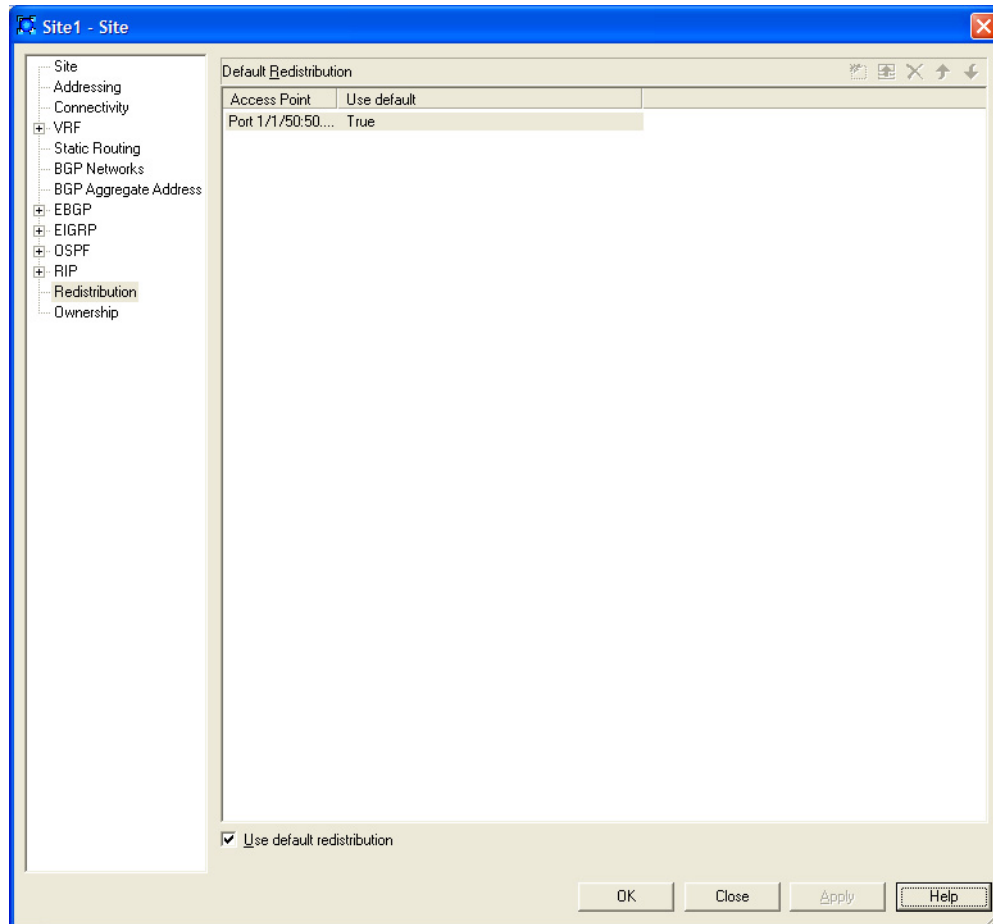
To avoid introducing routing loops and convergence problems, you can filter and refine the redistribution of routes by associating a manually pre-configured route map, or policy statement, with redistributed routes.

Directly-connected networks can also be redistributed into routing protocols. Service Activator supports direct redistribution of connected routes.

The default route may also be distributed via iBGP to peers within the VPN.

To specify metrics for route redistribution:

1. In the **Site properties** dialog box, select the **Redistribution** page.



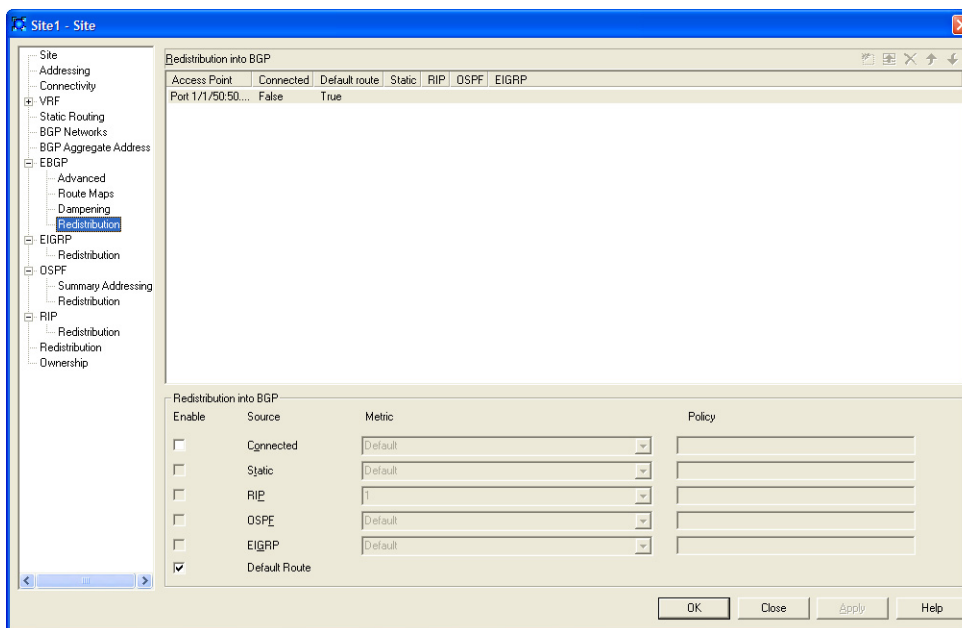
2. Clear the **Use default redistribution** checkbox.

The **Metric** and **Policy** fields that are enabled on the <Destination-protocol> Redistribution pages depend on the protocol selected for PE-CE connectivity.

**Note:** If the **Use default redistribution** checkbox is selected, the Service Activator default metrics will be used. Otherwise, the metric, and policy attributes for redistribution of connected routes can be specified in the redistribution matrix.

3. Select one of the following <Destination-protocol> Redistribution pages:
  - EBGP
  - EIGRP


- OSPF
  - RIP
4. If you wish to specify a metric and/or route map for connected routes, select the **Enable Connected** checkbox.
  5. Fill in the fields in each column for each connectivity type as required:  
For example, on the EBGp Redistribution page:



- **Metric:** The metric to apply to routes learned from the PE-CE protocol as they are redistributed into BGP.
- **Policy:** The name of a manually pre-configured route map or, for Juniper M-series devices, policy statement to apply to routes distributed into BGP.

On the other <Destination-protocol> Redistribution pages:

- **Protocol Metric:** The protocol-specific metric to apply to routes distributed from another protocol into the protocol used for PE-CE connectivity.
  - **Protocol Policy:** The name of a manually pre-configured route map or, for Juniper M-series devices, policy statement to apply to routes distributed from another protocol into the protocol used for PE-CE connectivity.
6. If you wish to distribute the default route via iBGP to peers, select the Enable Default Route checkbox.

7. To confirm your changes, click the Modify button.
8. Click the Set button to confirm your changes: 
9. Click **OK** to commit the changes and close the dialog box.

Hints and tips

- On the OSPF Redistribution page, and the RIP Redistribution page, and the EIGRP Redistribution page, only the value entered for BGP affects device configuration. A value specified for redistribution from any other protocol affects configuration only where two or more interfaces on the PE device participate in the same VPN, use different protocols for PE-CE connectivity and share the same VRF table.
- The RIP metric is based on hop count, and the maximum valid metric is 16. A value of 16 is considered infinite. We recommend you use a low metric when redistributing a protocol's routes into RIP.
- We recommend you apply a metric when redistributing connected routes.
- If no values are specified on a <Destination-protocol> Redistribution page, Service Activator applies default metrics to redistributed routes.

**Redistribution Matrix Columns**

As an example, with OSPF selected as the connectivity type, the column headings going across the redistribution matrix are interpreted as follows:

Metric	Type	Policy
Default	2	
Default	2	
Default	2	
Default	2	
0	2	

<b>Redist Matrix Columns</b>	<b>Column 1 Metric</b>	<b>Column 2 Type</b>	<b>Column 3 Policy</b>
<b>Meaning</b>	Redistribution of RIP into OSPF Metric value.	Redistribution of metric Type 1 from RIP into OSPF if selected. (Default is Type 2.)	Redistribution policy for RIP routes into OSPF.

## Setting up OSPF properties for a site

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### Cisco commands

The relevant Cisco commands for configuring OSPF on the PE router are configured automatically by Service Activator. They are as follows:

```
router ospf process-id vrf vrf
```

Configures OSPF in the context of the VRF table.

```
router-id id
```

This command is only configured if there is VRF reduction between two or more tables that use OSPF as the PE to CE protocol. The *id* is set to the IP address of the interface owning the VRF OSPF instance.

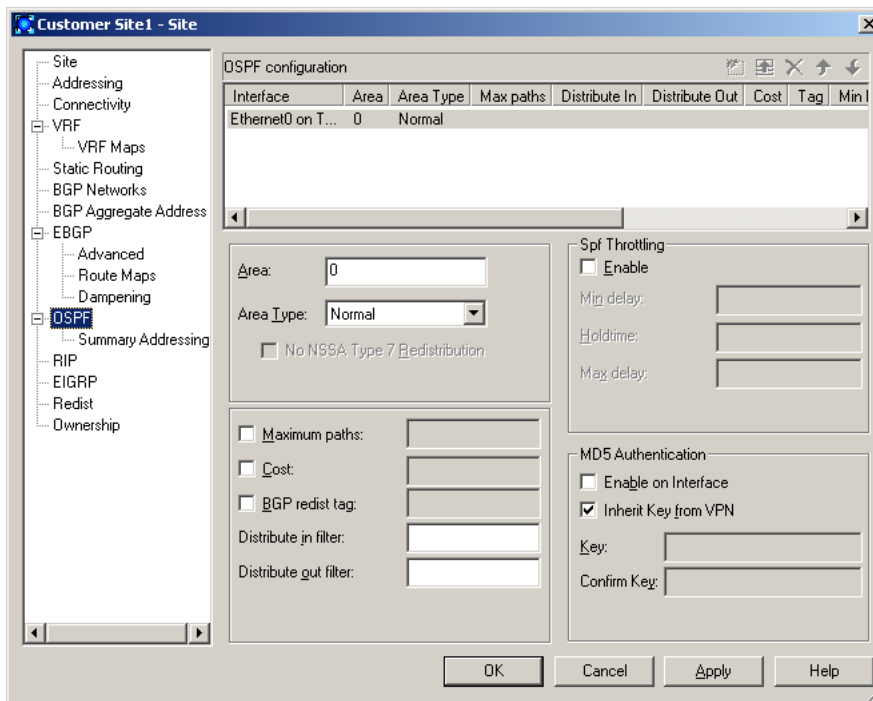
```
network ip-address mask area area-no
```

Specifies a network which identifies the interfaces that OSPF will run on. The PE-CE connection is always configured as area 0.

### To set up OSPF properties for a site

1. Display the **Site** dialog box and select the **OSPF** page.

- Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.



- Specify OSPF settings for each listed PE interface or sub-interface including **Area, Area Type, No NSSA Type 7 Redistribution, Maximum Paths, Cost, BGP redist tag, Distribute in filter, and Distribute out filter, under SPF Throttling, Enable, Min delay, Holdtime, and Max-delay, under MD5 Authentication, Inherit from VPN, Enable, Key, Confirm Key.**

There is additional details on OSPF Area Types and MD5 Authentication. For complete dialog box and property page descriptions, refer to the *Online Help*.

### Setting up OSPF Summary Addressing

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

#### To set up OSPF Summary Addressing for a site:

- Display the **Site** dialog box and select the **Summary Addressing** property page.



Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.

2. Specify values for each listed PE interface or sub-interface including **IP Address, Mask, Suppress Advertise, and Use this tag.**

## Setting up RIP properties for a site

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To set up a site's RIP properties

1. Display the **Site** dialog box and select the **RIP** page.

Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.

2. Configure the following for each listed PE interface or sub-interface:

- **Ignore Routes From**
- **Passive Interface**

There is detailed information available on these two panels. For complete dialog box and property page descriptions, refer to the *Online Help*.

## Setting up EIGRP properties for a site

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To set up a site's EIGRP properties

1. Display the **Site** dialog box and select the **EIGRP** page.

Any PE interfaces or sub-interfaces that you have linked to the site are listed. If no PE interfaces or sub-interfaces have been linked to the site, no addresses appear.

2. Configure the **MD5 Authentication** for each listed PE interface or sub-interface.

Note that the EIGRP ASN settings are configured on the **Site properties - Connectivity** page.

## MD5 Authentication

Use this panel to enable MD5 Authentication. Select the **Inherit from VPN** checkbox to inherit the interface's EIGRP authentication settings from the parent VPN. (The default settings for the VPN can be set on the **VPN properties - Connectivity page**.) Clear it for no authentication, or to allow selection of the Enable checkbox for MD5 Authentication.

EIGRP MD5 Authentication uses Key Chains which must be present on the device, either through previous manual configuration, or through a policy.

## Setting up the VPN

In order to set up a VPN, you need to create a VPN object and link the appropriate sites to it. You can also create a map to show the VPN and its connected sites.

## Applying QoS to CE devices or SAA to a VPN

If you need to apply QoS to CE devices, or SAA to a VPN, the order of setup tasks is significant and must be as follows:

1. Create a management VPN and propagate it to the network.  
The management VPN type provides control of the CE devices.
2. Create the fully-meshed or hub and spoke VPN and propagate it to the network.
3. Apply QoS or SAA to the fully-meshed or hub and spoke VPN.

For an outline of the steps required to set up a management VPN, see [Appendix A on page 145](#).

When removing a VPN to which QoS or SAA has been applied, you must remove the policy or measurement before deleting the VPN.

## Creating an MPLS VPN

VPNs are created for Service Activator customers – you cannot create an MPLS VPN that is customer-independent. For information on creating customers, see [Setting up customers on page 20](#).

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

### To create a VPN

1. On the **Service** tab, open the relevant customer folder and select the **VPNs** folder.
2. Right-click and select **Add VPN** from the pop-up menu.  
The **VPN** dialog box opens.
3. Specify values including **Name**, **Remarks**, **Level**, and **VPN Protocol**.
4. Select the **Connectivity** page and specify the **Connectivity Type**:
  - **Mesh**: Each site can communicate with all other sites.
  - **Hub and Spoke**: One or more sites act as a controlling interface.
  - **Management**: Hub and spoke topology that provides connectivity to the CE device where QoS or SAA will be implemented on the VPN.

If you are setting up a hub and spoke or a management VPN you can specify that at least one of the sites is defined as a hub. For more information, see [Specifying a hub site on page 52](#).

### Setting Route Target numbers

The import and export policies of a VRF table are defined by route target (RT) numbers. An import policy only allows iBGP routes whose RTs match the RTs of the import policy to be imported into the VRF table. An export policy specifies which RTs are attached to iBGP routes exported from the VRF.

By default, Service Activator automatically creates two RT numbers per VPN called **Default** and **Default+1**.

- The **Default** value is based on the domain's Autonomous System Number (ASN) and the unique object ID assigned to the VPN by Service Activator. By default, this value defines the import and export policies of all sites if the VPN is fully-meshed (**Mesh**), or the import and export policies of the hub site if the VPN is a hub and spoke or management VPN.
- The **Default+1** value is the **Default** value incremented by 1. By default, this value defines the import policy of all hub sites and the export policy of all spoke sites if the VPN is a hub and spoke or management VPN. This value is not used if the specified VPN connectivity is **Mesh**.

You can define additional RT numbers, and assign to each RT number any combination of import/export policy and site behavior.

Note that the default RT numbers created by Service Activator are in the format:

*ASN:Number*

If you wish to use RT numbers in the format:

*IPAddress:Number*

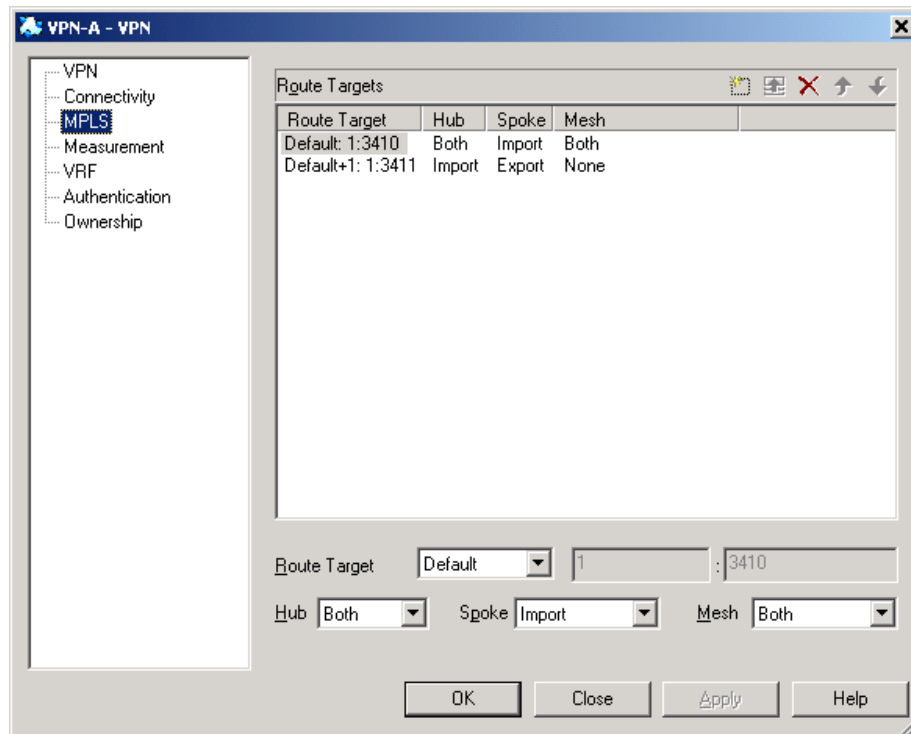
you must define custom RT numbers.

Service Activator checks that system-generated RT numbers are unique. However, no such check is made on user-defined RT numbers and non-unique numbers are permitted.

**Note:** For complete dialog box and property page descriptions, refer to the *Online Help*.

#### To define and allocate an RT number

1. Display the **VPN** dialog box and select the **MPLS** page.



2. Specify values including **Route Target** and **Hub, Spoke, Mesh**.

3. To save all the listed entries, select **Apply**.

## Linking sites to the VPN

You need to create the VPN by linking the appropriate customer sites to the VPN object. A site can be in more than one VPN and can be a hub in one VPN and a spoke in another.

### To link a site to a VPN

- Drag and drop the site object onto the VPN to create a link.  
If the VPN is a hub and spoke or management VPN, added sites are spokes by default.

Problems occur if spoke sites with separate VRF tables on a single PE device are added to a fully-meshed VPN while the device driver is down. The next time a transaction is committed after the driver has re-started the PE device is put into the 'Intervention Required' state and an error is raised. The problem does not occur if the VPN topology change is made after the device driver has re-started, however.

## Using an RD number per VPN or per site

By default, Service Activator automatically generates a site-specific VRF table name and RD number for each site that participates in a VPN.

At the VPN level, you can override the Service Activator default by specifying that the same VRF table name and RD number is applied to all sites that participate in

the VPN. You can choose whether to use Service Activator-generated values or specify your own VRF table name and/or RD number.

If you wish to use this feature, in addition to setting a VPN-level option you must also select the **Inherit from VPN** option in each relevant site's property pages. For more information, see [page 37](#).

Using a single RD number for all sites in a VPN is suitable only where a site belongs to one intranet VPN. If the site may become a member of an extranet VPN in the future, this method is not recommended.

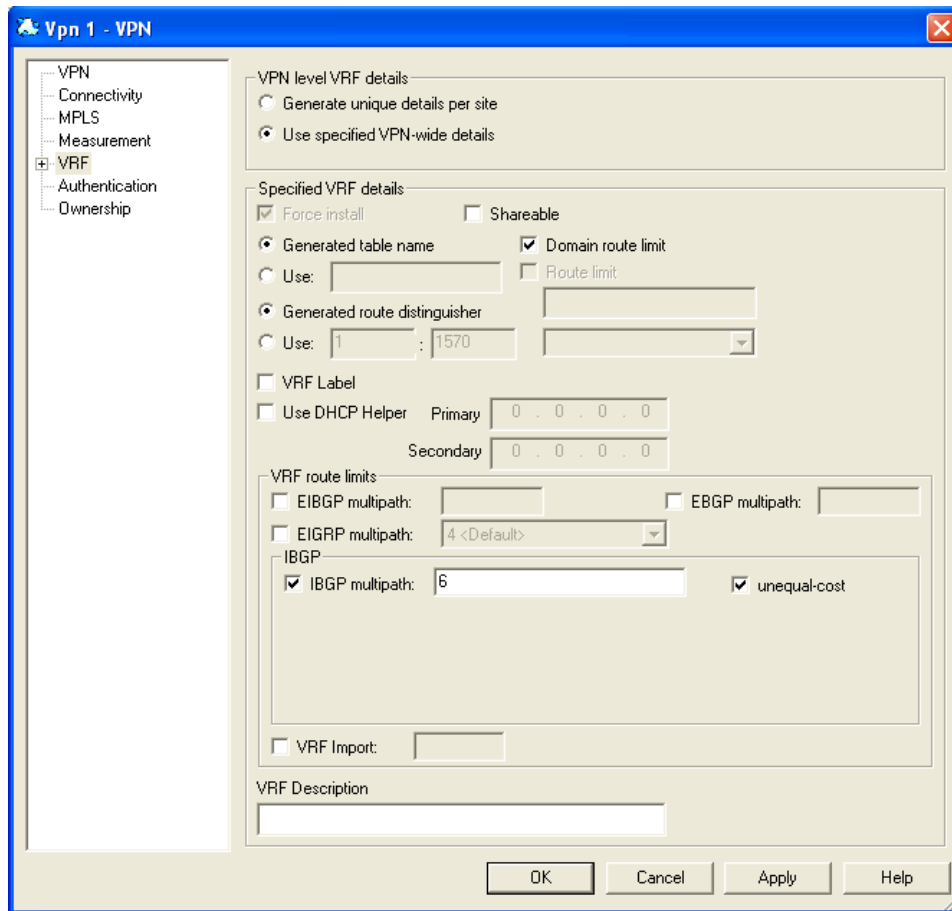
Use the **Multi-VPN override** option for more control over how RD numbers are assigned. If both **Inherit from VPN** and **Multi-VPN override are selected**:

- if the site is a member of only one VPN, the VRF table name and RD are derived from the parent VPN
- if the site is a member of multiple VPNs, the VRF table name and RD are derived using the site specific options

Note: **For complete dialog box and property page descriptions, refer to the Online Help.**

#### **To use the same RD number for all sites in a VPN**

1. Display the **VPN** dialog box and select the **VRF** page.



2. Select **Use specified VPN-wide details**.
3. Select values including **Force install**, **Shareable**, **Generated table name**, **Generated route distinguisher**, **Domain route limit**, **Route limit**, **VRF label**, **Use DHCP helper**, **EIBGP multipath**, **EBGP multipath**, **IBGP multipath**, **VRF Import** and **VRF Description**.
4. Enable the check box **unequal-cost** to allow unequal cost load balancing by selecting iBGP paths that do not have an equal cost.
5. The VRF description can also be set at the **Site** level. See [To set advanced VRF table options on page 37](#).

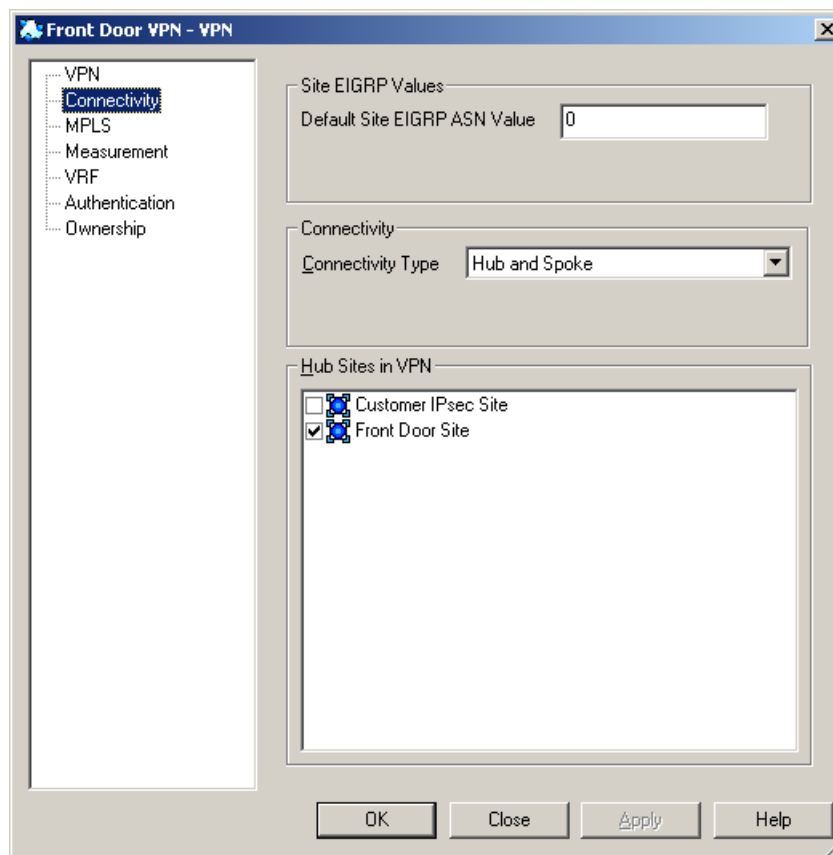
## Specifying a hub site

If you are setting up a hub and spoke or management VPN, you can specify one or more sites as the hub site.

### To specify that a site is a hub

1. From the VPN's pop-up menu, select **Properties** and select the **Connectivity** page.

Sites that have been added to the VPN are listed in the **Connectivity** page.



2. In the **Connectivity** field, select **Hub and Spoke** or **Management** from the drop-down menu.
3. To specify that a site is a hub, select the checkbox next to the relevant site's name.



## Creating a VPN map

Each VPN can be shown as a map view, which shows the sites within the VPN. You can choose whether the sites are laid out automatically or arrange them manually. In the manual layout option, you can specify whether objects snap to a grid layout and the granularity of the grid. The default is manual layout.

### To display the VPN map

- Double-click on the relevant VPN in the **Hierarchy** pane.

A representation of the VPN is shown in the **Details** pane. Initially it consists of the VPN object only. Sites linked to the VPN are listed in the palette.

### To specify layout options for a VPN map

1. If no map exists for this VPN, then from the VPN icon's pop-up menu, select **Add Map View**.
2. Add the map name. Optionally, you can also add a description, change the zoom level, and uncheck the Default Palette.
3. If a map exists for this VPN, then from the VPN map's pop-up menu, select **Properties**.

The **Map View** dialog box opens.

If a map has a background image, you may need to click on the map's tab at the bottom of the **Details** pane to display its pop-up menu.

4. Select the **Layout** page.
  - If you want to lay out sites manually, select **Manually lay out items on map**.
  - To specify that sites snap to a grid, select the **Snap to grid** option and specify the granularity of the **Grid** in millimetres.
  - If you want to lay out sites automatically, select **Automatically lay out items on map** and select items to be shown including **Networks, Sites, Devices, Interfaces, VCs, VC End Points**, and **Segments**.
  - You can specify **Max nodes**, **Max fan** and **Max Devices** values to control the automatic layout. For information on these settings, see *Network Discovery and Basic Setup*.

5. Click **OK**.

If you specified automatic layout, all previously unmapped sites are added to the VPN map.

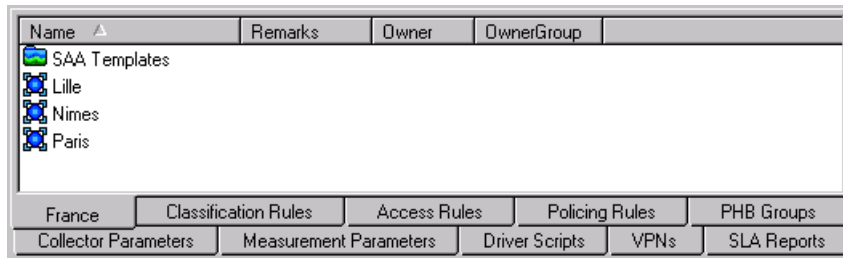
## Listing the sites in a VPN

You can list the sites that are associated with a VPN and display summary information for each site.

### To list the sites in a VPN

1. Select the relevant VPN from the hierarchy tree or the topology map.
2. Click the **Report View** button on the toolbar.

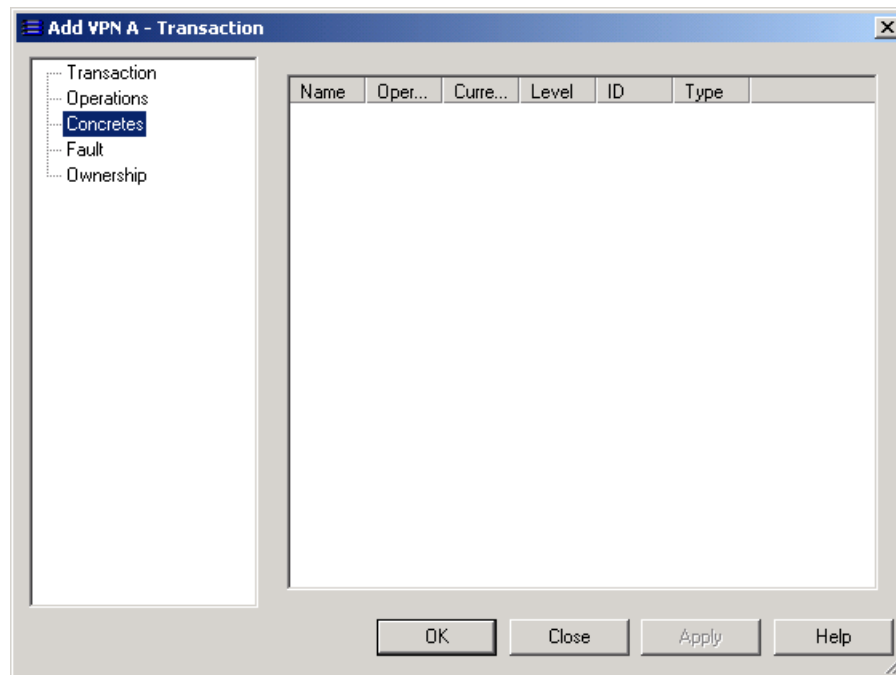
Service Activator lists the sites that are associated with the VPN and displays the properties for each site.



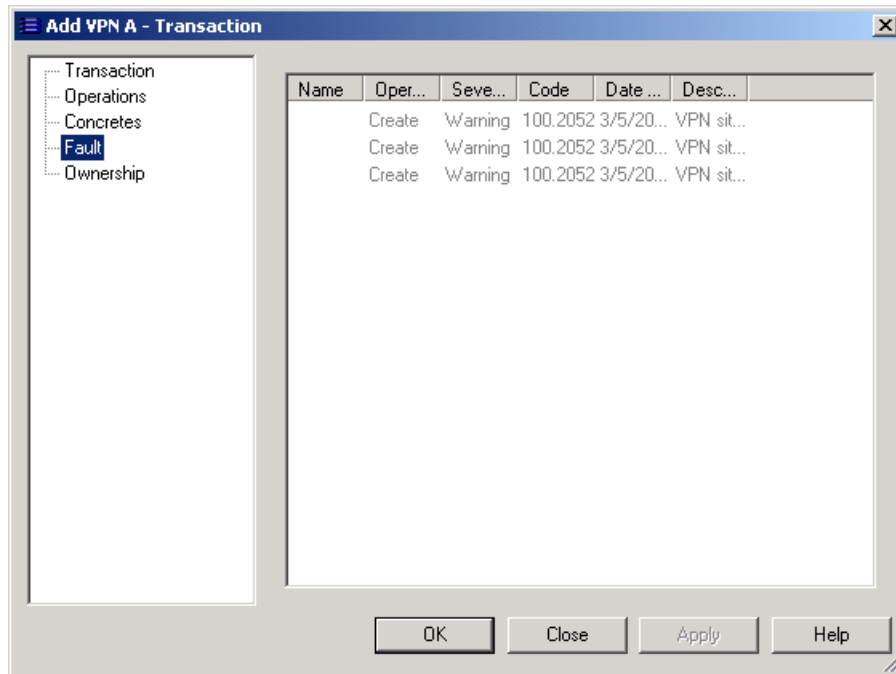
## Implementing the VPN

Once the site and VPN details are set up and the relevant devices are managed, the entire configuration can be applied by committing the transaction.

When you commit the transaction, any concrete VPNs that will be created are listed in the **Concretes** page of the Transaction dialog box.



Any validation errors are reported in the **Fault** page of the Transaction dialog box and the **Current Faults** pane.



If you wish to cancel the transaction after reviewing the concrete VPNs that will be created and the faults generated by the transaction, click **Cancel**.

If you wish to proceed with the transaction, click **OK**. Configuration details are sent to the proxy agent and on to the appropriate device drivers. For information on committing a transaction, see *Network Discovery and Basic Setup*.

After committing the transaction, you can check the configuration that has been applied to the routers by checking the device logs and by using Telnet. For more information, see the *Administrator's Guide*.

## Viewing implemented VPNs

You can view a list of the VPNs that have been propagated to the network and installed on an interface or subinterface.

### To view implemented VPN details

1. Double-click the relevant object (e.g. subinterface, interface or VPN) in the hierarchy tree or the topology map.
2. In the details pane, click on the **VPNs** tab to view VPNs implemented on the selected object. All concrete VPNs appear on a yellow background.

VPN/TLS	Site	Interface	Device	State	Conflict	ID
France	Nimes	at-1/3/0.0 on jm20-3		Inactive	False	1604
France	Lille	at-1/0/0.0 on jm10-1		Inactive	False	1600

VPN details are listed under the following headings:

- **VPN** – name of the VPN
- **Site** – the site associated with the VPN
- **Interface** – the interface associated with the site
- **Device** – this column remains blank for MPLS VPNs
- **State** – current state of the VPN:
  - **Inactive** – the VPN has been created but has not been propagated to the proxy agents
  - **Active** – the VPN has been propagated to the proxy agents
  - **Rejected** – the VPN configuration was rejected
  - **Installed** – VPN configuration has been installed on the designated device
- **Conflict** – there is a configuration error in the VPN
- **ID** – internal ID number by which the VPN is identified.



## Chapter 2

# Setting Up Layer 2 Martini VPNs

This chapter describes how to configure Layer 2 Martini VPNs. It includes the following major topics:

- a summary of the Layer 2 Martini VPN functionality on Cisco Devices
- pre-requisites for configuring Layer 2 VPNs
- procedures for creating, modifying, and deleting a Layer 2 Martini VPN
- a technical description of Layer 2 Martini VPNs

## Layer 2 Martini VPNs

A Layer 2 Martini point-to-point connection is a pseudo-wire (or tunnel) configured between two endpoints across an IP network.

The connection uses MPLS labels to encapsulate and transport various Layer 2 data formats, including VLAN to VLAN, Ethernet, Frame Relay, ATM Cell and ATM AAL5, across an IP network. The tunnel provides a transparent connection, so users see no change in their Layer 2 data. (Note that the tunnel does not aim to meet QoS aspects of the connection, particularly in the ATM case.) The Martini endpoints can be interfaces, sub-interfaces, or other endpoint identifiers (VCI/VPI on ATM, DLCI on Frame Relay, or VLAN ID on Ethernet).

A Layer 2 Martini VPN is an association of Layer 2 Martini point-to-point connections.

## Benefits of Layer 2 Martini VPNs

Layer 2 Martini VPNs enable the encapsulation and transport of legacy data types over IP networks. As service providers upgrade their network core, connections between legacy networks can be maintained. Customers needing traditional connectivity over a third-party network can be served using the same IP core network, regardless of the packet types they need to transport. Additionally, the tunnel saves the complexity of carrying the customers routes across the network.

Service Activator supports Layer 2 Martini VPNs across Juniper M-series and Cisco IOS devices.

Support of Ethernet technologies permits operators to use inexpensive Metro-Ethernet solutions in the Local Area Network, reducing the rollout cost of new networks.

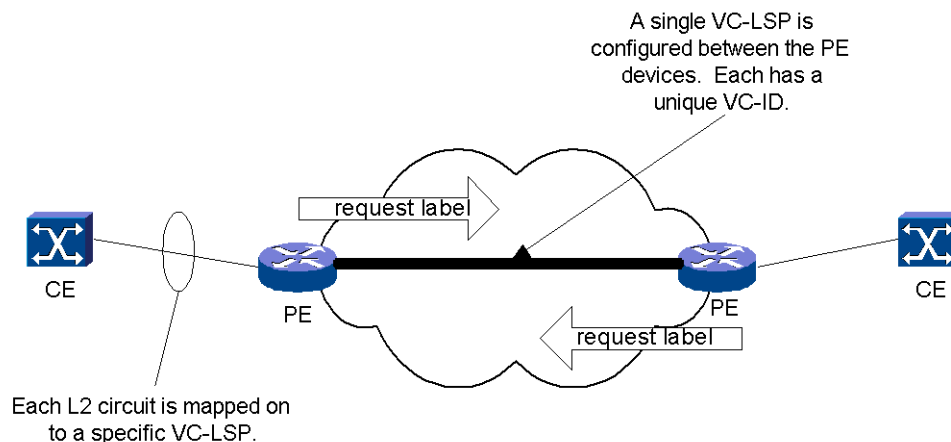
Similarly, Martini solutions can reduce the rollout costs associated with Mobile networks in transition from 2G to 3G. Using Martini tunnels, their 2G connection-oriented networks can traverse their new 3G IP core network. This saves the operational costs of supporting two different networks.

## Technical description of Layer 2 Martini VPNs

The Martini draft describes a mechanism that creates a bidirectional point-to-point connection between two PE routers. This connection is called a Pseudo-Wire (PW) or a virtual circuit - label-switched path (VC-LSP), and consists of two unidirectional LSPs.

Creating a PW (see [Figure 1](#)) is a 2-step process. The first step requires a Label Distribution Protocol (LDP) targeted-peering relationship to be established between the PE routers. Since routers can only exchange labels with their LDP peers, this step is required so that the typically non-adjacent PE routers can exchange labels.

**Figure 1: Martini point-to-point links**



The second step is to request a label for the layer-2 connection using an extension to the basic LDP signalling. A Layer-2 Forwarding Equivalence Class (FEC) is included in the label request that describes the circuit being connected together over the MPLS core. The returned label is mapped to the circuit described by the Layer 2 FEC and pushed on the bottom of the label-stack for each of the frames that



traverses the network. The PE at either end of the PW repeats this process of label requesting.

A PW is identified by a combination of its VC-ID and a Group-ID. The idea is that the Group-ID acts as a VC-LSP grouping mechanism (e.g. to identify a virtual interface value); the VC-ID then identifies the PW within the particular group. The Group ID can be seen to equate to the Layer 2 VPN identifier, though its use depends on the specific implementation.

Each Layer 2 frame or other data unit that arrives at the CE-facing interface of the PE is forwarded across the MPLS network with the label negotiated for the PW used to demultiplex it at the destination PE. Since each PE sent the label in response to the Layer 2 FEC, there is a unique mapping.

In addition to the PW-label, there is also a control word added to each incoming frame. This includes control flags and a sequence number, used to maintain the frame sequence in order-sensitive traffic.

The full processing of a frame includes:

- A frame arrives at the PE; its preamble and Frame Check Sequence (FCS) is removed.
- The control word is added to the front of the frame.
- The PW-label is added as the bottom-most label in the stack.
- The PE performs a lookup against the IP address for the destination PE and a transport label that reaches the destination PE is pushed onto the top of the stack - this may be either an RSVP tunnel or LDP LSP.
- The Transport label is swapped as the frame traverses the network and will be (in most cases) penultimate-hop-popped (removed) before reaching the destination PE.
- The destination PE will use the exposed PW label to determine the PW of the frame and from that determine the egress port - it is likely that will be a single step lookup in the label table.
- The control word may be used at this point to check the stream order and other relevant administrative functions.
- The FCS and preamble will be reformed and the frame transmitted on the CE-facing interface.

The exact processing varies according to the Layer 2 traffic type being supported.

A Layer 2 VPN can be constructed by associating a number of Pseudo-wires. The regular Martini service permits you to connect together PE devices in any arbitrary topology. In practice, when used as a service, Martini is frequently used to connect users to a server or other data centre in a hub and spoke topology. In the 3G wireless world, this hub and spoke topology is also employed since the logical

topology of such networks tends to fan out leaving a classic tree and branch graph topology with each vertex in the graph as a 3G network element.

There are fewer requirements to support a full-mesh topology due to the extra configuration the PE device needs to support in this case, however there are instances where this may be needed.

## Layer 2 Martini VPN devices and data types

This topic gives an overview of the different devices and data encapsulations supported by Service Activator in the configuration of Layer 2 Martini VPNs. It also gives specific details for VPN types in which there are variations from the typical configuration.

### Layer 2 Martini VPNs on Cisco routers and switches

Service Activator supports the configuration of Layer 2 Martini VPNs on Cisco IOS-based routers and switches which encapsulate and transmit a number of different types of data. The Martini endpoints can also be provisioned by Service Activator.

Cisco IOS-based equipment can be roughly categorized as either **switching IOS** or **non-switching IOS**.

**Switching IOS equipment:** Cisco equipment that supports Layer 2 (Ethernet) and Layer 3 (router) switching features, MAC learning, and VLAN bridging, typically in the Catalyst or Cisco 7600 range. Switching IOS equipment typically runs CatOS or Supervisor OS.

**Non-switching IOS equipment:** Cisco routers with none of the switching features described above. Devices in this category support standard IP routing between interfaces, run standard Cisco IOS, and are typified by equipment such as the 7200, 7500, 10700, and 12000.

### Layer 2 Martini VPNs on switching IOS Cisco devices

The following data types can be encapsulated on Layer 2 Martini VPNs on switching IOS Cisco devices:

Encapsulated data	Endpoints	Comments
Ethernet (Port)	Any combination of VLAN interfaces	Martini VLAN ID header is stripped on the Martini VC-LSP (Martini tunnel) and re-applied (if required) on the exit interface.
Ethernet (VLAN)	VLAN endpoints configured under Ethernet interfaces (not sub-interfaces)	See the notes below this table.

All Martini endpoints (such as DLCI, VLANs, VPI/VCI) and their parents (logical and physical interfaces) must have the role **Access** assigned.

For **Ethernet (Port)** encapsulation on switching IOS Cisco devices, a main interface is used. Endpoint VLAN IDs must be the same on both sides of the tunnel.

For **Ethernet (VLAN)** encapsulation on switching IOS Cisco devices, sub-interfaces are not used as the Layer 2 Martini VPN endpoints. You must create new or use existing VLAN endpoints. The endpoint VLAN IDs on both sides of the tunnel must be the same.

### Inter-operability between switching IOS and non-switching IOS devices

For inter-operability between switching IOS and non-switching IOS devices, VLAN mode (which retains the VLAN tag across the Martini VC-LSP) must be selected on the switched IOS devices. You must also connect to a VLAN VC identifier with the same VLAN ID.

## Layer 2 Martini VPNs on non-switching IOS Cisco devices

The following data types can be encapsulated on Layer 2 Martini VPNs on non-switching IOS Cisco devices.

Encapsulated data	Endpoints	Comments
Ethernet (Port)	Ethernet interfaces	All VLAN tags are preserved across the connection. Frames that enter the tunnel labelled VLAN n leave the tunnel labelled VLAN n.
Ethernet (VLAN)	VC identifiers	The VC identifier value represents the VLAN ID. The same VLAN ID must be used at both ends of the connection.
ATM Cell	Sub-interface with VC identifier	none
ATM AAL5	Sub-interface with VC identifier	none
Frame Relay	Main interface with VC identifier	The VC identifier value attached to the main interface must be created manually.

All Martini endpoints (such as DLCI, VLANs, VPI/VCI) and their parents (logical and physical interfaces) must have the role **Access** assigned.

ATM Cell Layer 2 Martini tunnel endpoints must have the same VPI/VCI. ATM AAL5 tunnel endpoints are not required to have the same VPI/VCI.

For Ethernet VLAN on non-switching IOS Cisco devices, VC identifiers are used to represent the VLAN ID. VC identifiers are configured on Ethernet sub-interfaces and used as the Layer 2 Martini VPN endpoints.

For Frame Relay encapsulation on non-switching IOS Cisco devices, sub-interfaces are not used as the Layer 2 Martini VPN endpoints. You must manually pre-configure or use existing PVCs (Permanent Virtual Circuits) on the main interface.

## Layer 2 Martini VPNs on Juniper M-series devices

Service Activator supports the configuration of Layer 2 Martini VPNs on Juniper M-series devices which encapsulate and transmit a number of different types of data.

The following data types can be encapsulated on Layer 2 Martini VPNs on Juniper M-series devices.

Encapsulated data	Endpoints
Ethernet (Port-based)	Main interface
Ethernet (VLAN-based)	Sub-interface with VC identifier
ATM Cell	Sub-interface with VC identifier
ATM AAL5	Sub-interface with VC identifier
Frame Relay	Sub-interface with VC identifier

All Martini endpoints (such as DLCI, VLANs, VPI/VCI) and their parents (logical and physical interfaces) must have the role **Access** assigned.

ATM Cell Relay Layer 2 Martini tunnel endpoints must have the same VPI/VCI. ATM AAL5 tunnel endpoints are not required to have the same VPI/VCI.

Note: When creating a Layer 2 Martini VPN with SONET interfaces on Juniper M-series devices as endpoints, the MTU values must match. Note that this must be set manually. Service Activator does not validate the MTU values, so you will not be notified when there is a potential mismatch. The Martini circuit will be created in the GUI but may not be operational if the MTU values do not match on the SONET endpoints.

### Martini Layer 2 VPNs on Huawei devices

Service Activator supports the configuration of Martini Layer 2 VPNs on Huawei devices. Ensure that Huawei devices are configured to support Martini Layer 2 VPNs.

Refer to the Huawei Operation Manual for information.

### Overview of Layer 2 Martini VPN creation

This section summarizes the activities involved in creating Layer 2 Martini VPNs.

#### Pre-requisites for configuring a Layer 2 Martini VPN

- Discover devices and assign roles
- Create customers
- Check interface capabilities
- Discover and pre-configure all sub-interfaces
- Complete other pre-configuration requirements

### Create the Layer 2 Martini VPN

- Add the Layer 2 Martini connections.
- Set the options in the property page.
- Assign the endpoints to the new Layer 2 Martini tunnel.

## Discovering devices and assigning roles for VPN setup

When you have set domain-level information, you can run the discovery process to find all the P and PE routers in the network and include their details in Service Activator's database.

All devices within the network must be correctly assigned system-defined roles, that is, PE routers must be classified as gateway devices, P routers classified as core devices and CE routers, if visible, classified as access devices. The recommended way of assigning roles is by means of role assignment rules, which automatically assign roles during device discovery.

All interfaces within the network must be correctly assigned system-defined roles:

- On CE (access) devices, the interface connected to the PE device must be classified as an access interface. Interfaces connected to local segments must be classified as local interfaces.
- On PE (gateway) devices, the interface connected to the CE device must be classified as an access interface. Interfaces connected to other PE devices or P (core) devices must be classified as core interfaces.
- All interfaces on P (core) devices should be classified as core interfaces.

### To discover the network

1. Choose **Discover** from the **Discovery** menu.  
The Topology Discovery dialog box is displayed.
2. On the **Discovery** page, enter the DNS name of the IP address of each device to be discovered.  
Optionally, set the **Hops** field to a value between 1 and 10.
3. Press **OK**. You are prompted to save the changes by choosing **Save** from the **File** menu. As soon as the changes are committed to the database, the device discovery process starts.

### Hints and tips

- The **Discovery** menu option is not available if there are unsaved changes in the user interface. You must either commit or save the current transaction before you can run a discovery.

- In an MPLS domain, the core provider network is assumed to use public addresses, and the hop count can be used within the core network. All CE routers are assumed to use private addresses and an IP address or DNS name must be specified in order to discover them.
- You may need to change the default settings on the **SNMP** page.

### To assign roles to devices and interfaces

All devices within the network must be correctly assigned roles (i.e. PE routers classified as gateway devices, P routers classified as core devices and CE routers, if visible, classified as access devices.)

This will be done automatically if you have set up role assignment rules, otherwise you need to manually assign a role to each device and interface to be managed.

#### Hints and tips

- You are advised to set up role assignment rules to classify devices and interfaces correctly.

### To manage a device

Before a device can be managed by Service Activator, you also need to ensure the following:

- All devices in the domain that are to be managed by Service Activator must be assigned to a proxy agent. Although it is possible to assign devices manually, it is generally performed automatically during device discovery.
- All devices to be configured by Service Activator need to be set to Managed. When devices are first discovered, their status is set to Unmanaged. To set all devices to Managed, select the network and choose **Manage All Devices** from the pop-up menu.

## Creating a customer

You must create a customer before you can create a VPN.

### To set up a customer

1. Choose the **Customers** folder on the **Service** tab in an explorer window and choose **Add Customer** from the pop-up menu.

The **Customer** dialog box is displayed.

2. Enter the following:
  - **Customer name:** Specify an identifying Name for the customer.
  - **Remarks:** Additional comments (optional).
  - **Reference:** Customer reference number (optional).

3. Press **OK** to close the dialog box.

## Checking Interface Capabilities

Before creating a Layer 2 Martini VPN, check the capabilities of the interfaces, sub-interfaces or provisioned sub-interfaces on the devices. You need to determine if they will support the endpoints for the Martini tunnel.

### To check the interface capabilities for supporting a Layer 2 Martini VPN

1. Right click on the interface and select **Properties**.
2. Display the **Capabilities** property page.
3. Under **Outbound Properties**, expand **Interface Creation Support**.
4. Ensure that the type of encapsulation you wish to use in your Layer 2 Martini VPN is supported by the interface.
5. Confirm that the **Role** for the interface is set to Access.

## Completing other pre-configuration for Layer 2 Martini VPNs

Ensure that devices are pre-configured as described in this section, before configuring the Layer 2 Martini VPN.

- MPLS must be enabled on all appropriate interfaces.
  - `mpls label protocol ldp` - specifies the use of the LDP label distribution protocol
  - `tag-switching ldp router-id Loopback0 force` - enable tdp tag-switching, force the Loopback0 address to be used as the router ID
- The `ip cef` or `ip cef distributed` command must be manually configured in order to turn on CEF or dCEF.
- On PE devices, an IGP such as OSPF or EIGRP must be configured in order to distribute IP routes. These are required for IP connectivity, and to enable labels to be allocated by the separate LDP (Label Distribution Protocol) or TDP (Tag Distribution Protocol).
- Tag-switching of IPv4 packets on the WAN-facing (Core-facing) interfaces. (These are not the same interfaces on which sub-interfaces for the Layer 2 Martini VPN tunnel endpoints are to be configured.)
  - `interface <interface name>` - specify WAN facing interface for next command.
  - `tag-switching ip` - enables tag-switching of IPv4 packets on the specified interface and device



- Devices used in Layer 2 Martini VPNs should be configured to use the Gateway role. Interfaces used as endpoints should be configured to use the Access role.
- Ensure that a loopback interface exists on Cisco devices on which you are configuring Layer 2 Martini services.

If a no loopback interface is present, you will receive a generic configuration error. Note that the configuration will continue to be resent to the device until you create a loopback interface.

### **Manual pre-configuration: specify LDP protocol on interfaces for Martini L2 connections**

Specify the Label Distribution Protocol on each interface to be used for a Layer 2 Martini connection. If you do not specify LDP, tag distribution protocol (TDP) is used instead.

Log into the PE router and enter: `mpls label protocol ldp`

### **Manual pre-configuration: assign LDP Router IDs to the PE Routers**

To assign LDP router IDs to the PE routers, perform the following steps. Both PE routers require a loopback address that you can use to create a virtual circuit between the routers.

1. Enter interface configuration mode: `interface loopback0`

Note: The LDP router ID must be configured with a 32-bit mask to ensure proper operation of MPLS forwarding between PE routers.

2. Assign an IP address to the loopback interface: `ip address <ip-address>`
3. Assign the loopback IP address as the router ID:

```
mpls ldp router-id loopback0 force
```

Note: This command forces the loopback interface to be the LDP router ID on each PE router. Without "force", the router can assign a different router ID, thereby preventing the establishment of Virtual Circuits between PE routers.

### **Manual pre-configuration: Martini circuits on Ethernet interfaces**

If any physical interface encapsulation incompatibilities pre-exist on the router, Service Activator detects them when the device driver is building a new configuration for Martini circuits. An error is displayed in the UI, and you are given the option to manually correct the interface encapsulation.

In order to expedite the configuration process, ensure that the following manual configuration exists on the router:

1. For 802.1Q VLANs and/or VLAN-based L2circuits ensure that vlan-tagging is enabled on physical interfaces and that each logical subinterface has a VLAN ID configured.
2. For physical Ethernet interfaces to be used in port-based Martini circuits, ensure that there is no vlan-tagging and either only unit 0 or none of logical subinterfaces are present.

For more information, see the *JUNOS Internet Software MPLS Applications Configuration Guide*.

### Huawei device pre-configuration

Ensure that Huawei devices are configured to support Martini Layer 2 VPNs.

Refer to the Huawei Operation Manual for more information regarding basic MPLS core, SNMP and Telnet configuration.

### Provisioning endpoints (VC IDs) for a Martini L2 connection

If you are provisioning a Layer 2 Martini VPN encapsulating Frame Relay, Ethernet or VLAN data on Cisco equipment, sub-interfaces are not used.

The following Layer 2 Martini connection requires VC identifiers as endpoints when encapsulating:

- Frame Relay (on non-switching IOS Cisco equipment)

The following Layer 2 Martini connections require interfaces with configured VLAN IDs as endpoints when encapsulating:

- Ethernet (on switching IOS Cisco equipment)
- Ethernet VLAN (on switching IOS Cisco equipment)

**Note:** The endpoints for Layer 2 Martini VPN encapsulating Ethernet or VLAN on switching Cisco equipment must be configured manually outside of the Service Activator UI. In addition, modification or removal of these endpoints cannot be performed inside Service Activator - you must modify or remove them manually.

**Note:** Refer to [Layer 2 Martini VPN devices and data types on page 62](#) for details about the different devices and data encapsulations supported by Service Activator for Layer 2 Martini VPNS, the Martini endpoints required, and details about VPN types for which there are variations from the typical configuration.

**To create a VC identifier for a Layer 2 Martini VPN encapsulating Frame Relay data on non-switching Cisco devices:**

1. In the hierarchical tree, expand the device containing the interface on which you are provisioning the PVC. Alternatively, double click the device and the interface in the **Topology** map.
2. Double click the interface to display the **Details** window.
3. Log into the device and configure the following commands manually on the device:

```
interface <interface>
encapsulation frame-relay ietf
frame-relay intf-type dce
exit
connect fr1 <interface> <PVC identifier> l2transport
exit
```

4. Rediscover the device in the Service Activator GUI.

**To create a VLAN endpoint for a Layer 2 Martini VPN encapsulating Ethernet (port) or Ethernet VLAN data on switching Cisco devices:**

Configure the VLAN endpoint manually on the device. The required configuration is as follows:

```
Ethernet (port):
interface <ethernet port>
encapsulation dot1Q <vlan id>
exit
```

**Note:** When a VLAN endpoint is not associated with a Martini point-to-point VPN object, it is in 'Conflict' state. When the endpoint is attached to the point-to-point object its state can reflect its status.

**Creating a Layer 2 Martini VPN**

In order to create a Layer 2 Martini VPN, you require Martini endpoints. Sub-interfaces can be created manually or by using the Service Activator subinterface creation feature. You must manually pre-configure any FR VC interfaces.

The Martini endpoints must be provisioned with the correct encapsulation for the type of Layer 2 Martini VPN you are creating. Then create the Layer 2 Martini VPN object itself, set the appropriate options, and finally, assign the relevant endpoints to the VPN.

Refer to [Overview of Layer 2 Martini VPN creation on page 65](#) for an overview of pre-requisite tasks.

**To create a Layer 2 Martini VPN:**

1. On the **Service** tab, open the relevant customer folder, select the **Point-to-Points** folder, right click, and select **Add L2 Martini-Pt-Pt** from the drop-down menu.

The **L2 Martini Pt-Pt** dialog opens.

2. On the **L2 Martini Pt-Pt** page:
  - **Name:** specify a name for the Layer 2 Martini VPN. The name may contain alphanumeric characters only, and may not include spaces.
  - **Remarks:** add any additional remarks (optional)
  - **Type:** choose the appropriate encapsulation type, matching the encapsulation selected when you pre-configured the sub-interface endpoints
    - ATM AAL5
    - ATM Cell
    - Ethernet
    - Ethernet VLAN
    - Frame
  - **Martini VC ID:** if **Automatic** is checked, Service Activator provides a VC ID for you. Otherwise, leave it unchecked and specify a VC ID.
3. If you wish to restrict access to the Layer 2 Martini VPN object, select the **Ownership** page and specify the details.
4. Add the Martini endpoints (interfaces, sub-interfaces, provisioned sub-interfaces or VC interfaces) to the Layer 2 Martini VPN by dragging the desired Martini endpoint objects into the new Layer 2 Martini VPN object. This selects them as the Martini endpoints.

For details on the **L2 Martini Pt-Pt** dialog box, refer to the **L2 Martini Properties** page in Online Help.

**Modifying Layer 2 Martini VPNs**

You can modify the attributes of an existing Layer 2 Martini VPN or attributes of the tunnel endpoints and the appropriate configuration changes will be made on the devices involved.

**To modify the properties of an existing Layer 2 Martini VPN**

1. In the hierarchical tree, select the **Service** tab, and expand the relevant customer folder.

2. Expand the **Point to Points** folder and locate the Layer 2 Martini VPN to be modified in the hierarchy, or double-click the **Point to Points** folder and locate the Layer 2 Martini VPN in the **Details** window.
3. Right-click the Layer 2 Martini VPN and select **Properties**.  
The **L2 Martini Pt-Pt** dialog box is displayed.

4. Make changes to one or more of the following fields:

- **Name**
- **Remarks**
- **Martini VC ID**

For details, refer to the **L2 Martini Pt-Pt** property page in the Online Help.

You can also make changes on the **Ownership** property page. For details, refer to the **L2 Martini Pt-Pt (Ownership Page)** in Online Help.

5. Click **OK**, and commit your changes.



## Chapter 3

# Setting Up Transparent LAN Services

This chapter describes how to configure a Transparent LAN Service (TLS) on Riverstone devices, and transparent VLAN service on Cisco and Extreme devices. The chapter:

- Provides a brief overview of TLS, identifying key concepts
- Describes the points you need to consider when planning a TLS
- Describes the preparation tasks associated with setting up a TLS, including manual pre-configuration, discovering the network and assigning roles, and setting up customers
- Explains how to create a TLS, and create and link the layer 2 sites that participate in the TLS
- Describes how to apply rate limiting to a site
- Describes how to implement a TLS and view details of the TLSs implemented on a port or device

## Overview

A TLS connects separate customer Ethernet LAN segments via an MPLS network. The connection across the network appears to the customer as a single LAN segment.

Service Activator supports the encapsulation and transport of layer 2 frames across the TLS as described in the Lasserre TLS Draft.

Configuration of the TLS occurs at Provider Edge (PE) devices within an MPLS network. Ethernet frames are mapped to a particular service instance based on a combination of the port on which they arrive at the PE device and, optionally, the 802.1Q tag that has been applied to them. As with other VPN solutions, inner and outer tunnels are used:

- Outer tunnels provide a transport mechanism between the PE routers in the TLS
- Inner tunnels, referred to as VC-LSPs, form a full mesh between the PEs in each TLS instance and are particular to that TLS.

Multiple VC-LSPs may be carried by a single transport 'outer' LSP.

### About core VLANs

TLS objects are used to connect to core VLANs (including stacked VLANs) created with the separately licensed **VLAN Activation Module**. For complete information on the VLAN Activation Module and on how to connect to a core VLAN through a TLS, **refer to the Service Activator on-line help**.

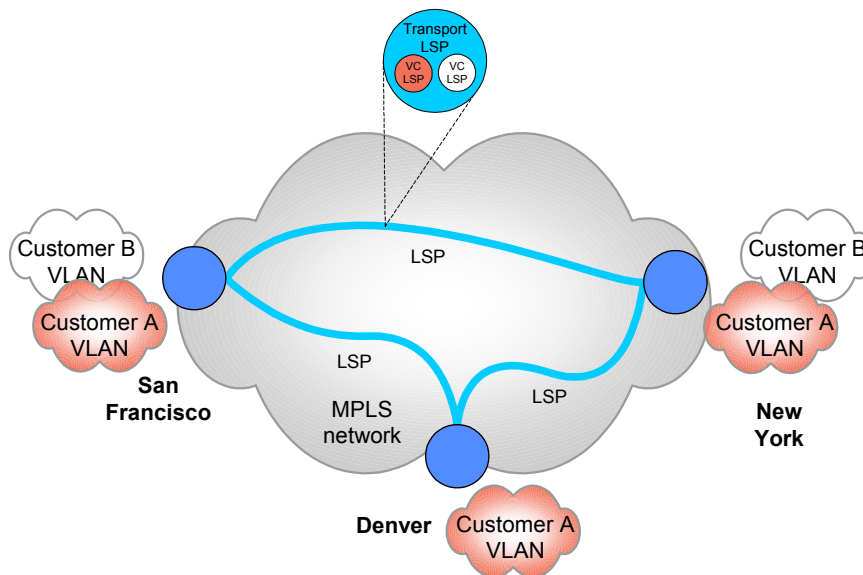
### VC-LSPs

The Lasserre TLS solution uses VC-LSPs as defined in the Layer 2 Martini over MPLS Internet drafts. A targeted LDP peering association between two PE devices creates the VC-LSP. The devices exchange information about the Layer 2 protocol that will be carried – in the TLS case, this is either untagged or 802.1Q tagged frames. This exchange also includes information about the TLS instance of which the VC-LSP forms a part. The Forwarding Equivalence Class (FEC) thus describes Layer 2 information, rather than the more usual IP prefix.

Each PE sends back a VC-LSP label, which is mapped to the FEC. When a frame is received at the PE, it examines its forwarding table and applies the correct VC-LSP label. The correct transport label is then added and the frame is forwarded to the correct destination. At the egress PE router, the VC-LSP label is used to identify the correct Ethernet port over which to forward the enclosed frame.

In the following diagram, VC-LSPs are configured for the Customer A TLS instance between San Francisco, Denver and New York. The VC-LSPs are contained within the transport LSPs that connect these destinations.





## Transport LSPs

Transport LSPs are responsible for linking PE routers together. Each VC-LSP must be forwarded to the correct PE by the transport LSP.

## 802.1 support

The IEEE standard 802.1Q describes a VLAN tag that can be applied to an Ethernet frame. The tag value is the VLAN ID, a number assigned to switches in an Ethernet network. Tagged frames can only be forwarded to switches that are configured with the same VLAN ID as the tag. Switches may be in more than one VLAN at a time, connected by trunk ports over which tagged frames are sent. Access ports to the Ethernet network may only be assigned a single VLAN ID. The frames arriving at an access port are untagged.

The 802.1P standard is not supported on Service Activator.

## Mapping Frames to the TLS

To complete the TLS service, a mapping must be established between incoming Ethernet frames to the PE and the VC-LSPs that are configured over the MPLS core. This mapping can be:

- Port based – all frames from a particular port are mapped to the service.

- VLAN based – all frames with an 802.1Q tag of a given value are mapped to the service.
- Port and VLAN based – all frames from a particular port with a given 802.1Q tag are mapped to the service.

Service Activator supports port-based, and port and VLAN-based TLSs. The mapping to the TLS instance is configured on the PE device.

On Riverstone devices, the TLS instance is identified by a customer profile and a customer ID. All the VC-LSPs with the same customer ID at the PE form the possible destinations for the Ethernet frames mapped to that TLS instance.

Ethernet is a broadcast service and this must be replicated in the TLS. Therefore, when a frame is received at a PE device for an unknown destination, it is forwarded over all the VC-LSPs in the TLS. When a response to this frame is received at the PE device, the device first learns which VC-LSP the frames were returned over before forwarding the frame over the correct Ethernet port. Future frames to that destination are then only sent over the learned VC-LSP. This mechanism is called 'flood and prune'.

A port that handles incoming traffic to the TLS may therefore receive tagged or untagged frames and tagged frames may belong to one or more VLANs. On Riverstone devices, ingress ports to the TLS may be configured as one of the following, depending on whether tagged or untagged frames are handled:

- A trunk port – receives and transmits tagged frames belonging to one or more VLANs; a trunk port may also be configured to transmit untagged frames by making it part of the native VLAN (typically VLAN 1), but this is not supported by Service Activator.
- An access port – receives and transmits untagged frames and frames belonging to a maximum of one VLAN. By default, access ports are considered to be part of the native VLAN (typically VLAN 1) unless they are explicitly assigned to another VLAN.

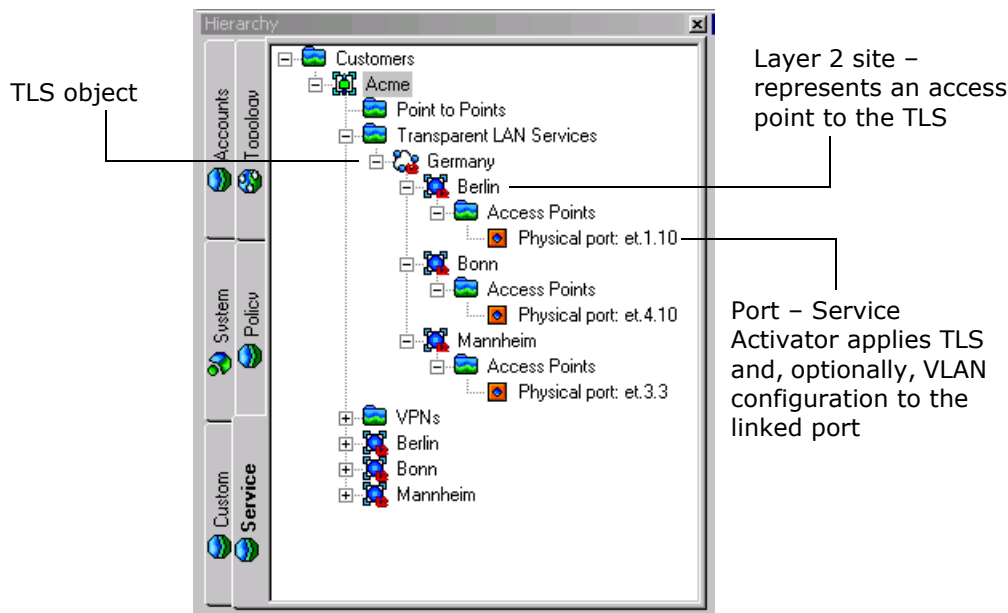
## Planning the TLS

Service Activator supports the following TLS types:

- Port-based – access to the TLS is controlled by incoming port number
- Port and VLAN-based – access to the TLS is controlled by incoming port number and VLAN ID. Incoming frames may already be tagged or a tag may be applied to them by Service Activator.

A TLS is represented by a TLS object in the user interface, and the edge points of the TLS are represented by layer 2 site objects. Each layer 2 site is linked to one or more ports that indicate where Service Activator will apply TLS configuration. If Service Activator is to tag incoming frames, VLAN configuration will also be applied.

A layer 2 site may include ports on both the PE and, optionally, the CE device.



Service Activator represents a layer 2 port as an interface object in the user interface. In the descriptions that follow, the term 'interface' is used when referring to TLS setup through the user interface.

Service Activator applies the concept of port and port and VLAN-based entry criteria both to the TLS object and the layer 2 sites that are linked to it:

- A port-based TLS consists of a number of port-based layer 2 sites
- A port and VLAN-based TLS consists of a number of port and VLAN-based layer 2 sites

## Creating a TLS

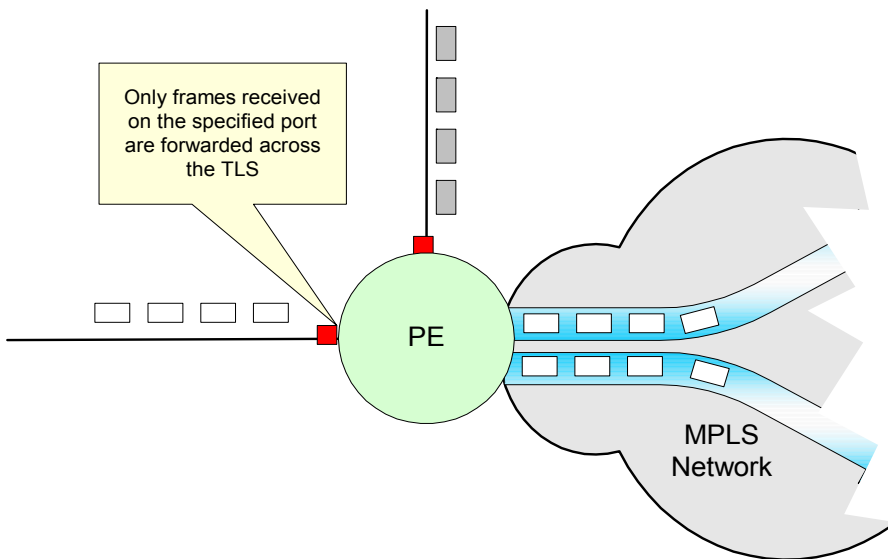
You create a TLS by:

1. Discovering the network and assigning the correct roles to the devices and ports that will participate in the TLS, described on [page 89](#).
2. Creating a TLS object, described on [page 91](#).
3. Creating layer 2 sites that represent the edge points of the TLS and linking them to the relevant ports, described on [page 94](#).
4. Linking the layer 2 sites to the TLS object, described on [page 96](#).

When you link layer 2 sites to a TLS, Service Activator configures a full mesh of LSPs between peer PE devices.

## Port-based TLS

In a port-based TLS, forwarding of frames across the TLS is based on incoming port number.



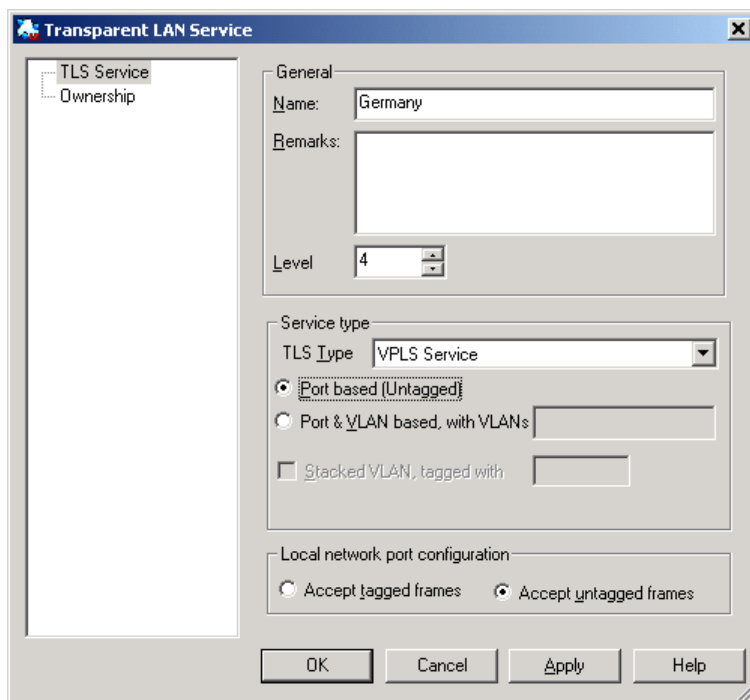
Incoming frames may be tagged frames or untagged frames.

Any VLAN configuration and management is performed by the service provider’s customer. Service Activator simply configures the specified ports at the edge of the TLS to be either Ethernet access ports or 802.1q trunk ports, depending on whether untagged or tagged frames are transmitted across the TLS.

You cannot perform tagging of incoming frames at a port-based site.

You specify whether the TLS accepts tagged or untagged frames when creating the TLS object.

**Note:** To create a port-based service, select **VPLS Service** in the **TLS Type** dropdown.

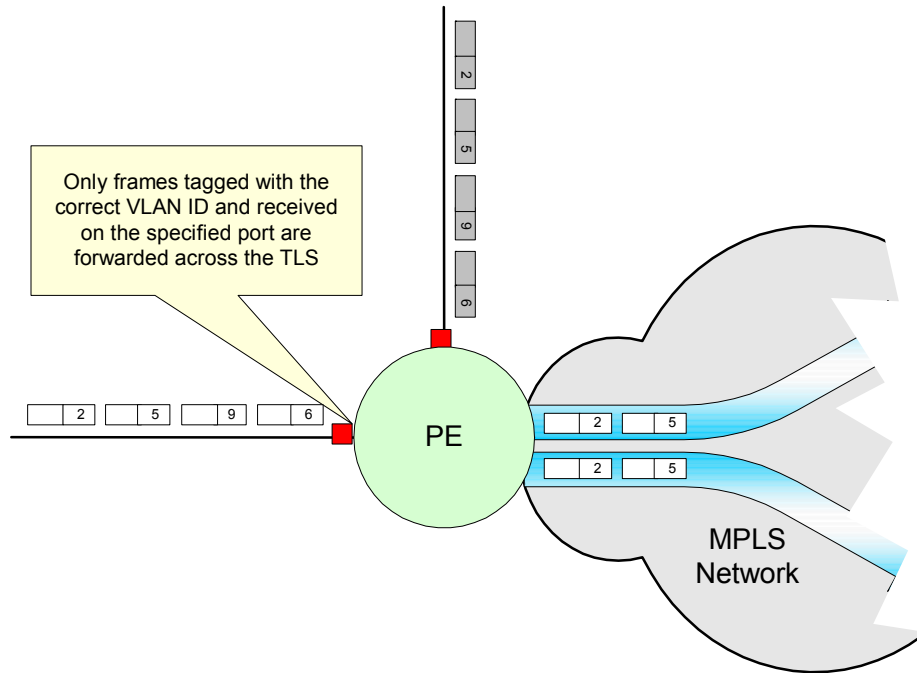


The edge points for the TLS are defined by port-based layer 2 sites. Each site may contain a single port. A port-based site may be linked to one port-based TLS.

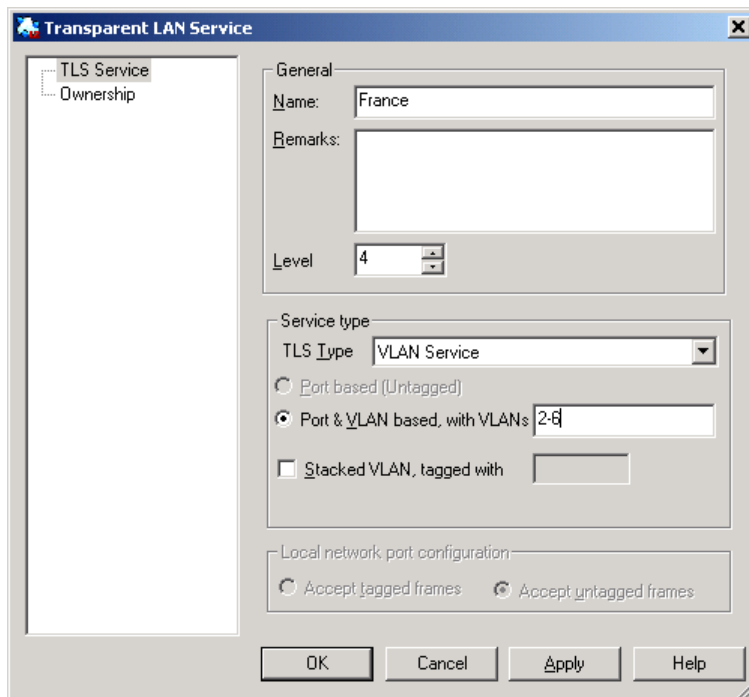
You specify on which ports incoming frames for the TLS will be received by linking the Access interface on the appropriate Gateway (PE) device to a layer 2 site. An interface can be linked to one port-based layer 2 site.

### Port and VLAN-based TLS

In a port and VLAN-based TLS, frame forwarding is based on incoming port number and the ID of the VLAN to which the frame belongs.



Incoming frames may have been tagged by the customer before reaching the entry point to the TLS, or you can specify that Service Activator tags incoming frames before forwarding across the TLS.



If Service Activator tags incoming frames, in addition to configuring the TLS it also configures the relevant VLANs and assigns them to the relevant ports. Frames are tagged at the PE device or, if the service provider has visibility of the CE device, tagging is performed here.

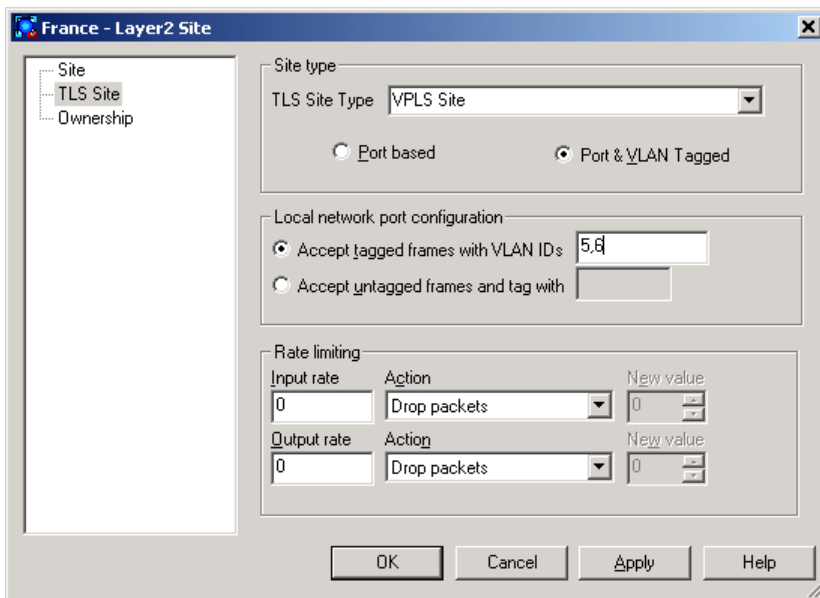
You specify which VLAN IDs will be carried by the TLS when creating the TLS object. These VLAN IDs form part of the customer profile that is configured on the PE device and specify the range of VLAN traffic the TLS will carry.

**Note:** If creating a TLS service on Riverstone equipment, select **VPLS Service** in the **TLS Type** dropdown. If creating a TLS to connect to a core VLAN configured with the **VLAN Activation Module**, select **VLAN Service** in the **TLS Type** dropdown.

Each layer 2 site may accept frames tagged with any of the VLAN IDs specified in the TLS definition. The VLAN IDs specified in the layer 2 site definition are used to create the necessary VLAN definitions on the incoming port.

**Note:** If creating a TLS service on Riverstone equipment, select **VPLS Service** in the **TLS Site Type** dropdown. If creating a TLS to connect to a core VLAN

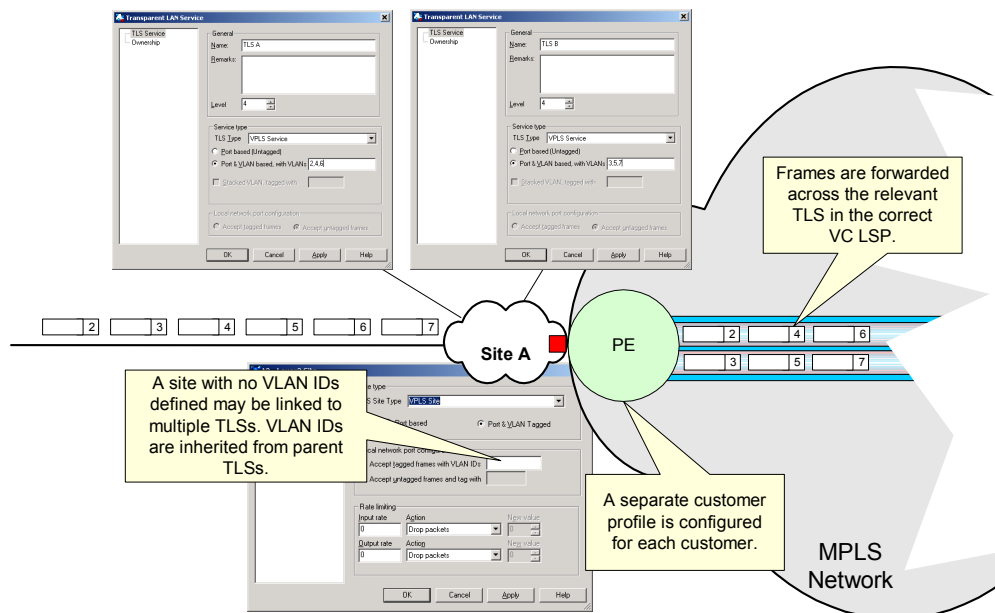
configured with the **VLAN Activation Module**, select **VLAN Service** in the **TLS Site Type** dropdown.



You may also choose not to specify VLAN IDs in the layer 2 site definition. If no VLAN IDs are specified for a layer 2 site, the VLAN IDs specified for the TLS are inherited to the site.

A layer 2 site that has no VLANs specified in its definition may be linked to multiple port and VLAN-based TLS objects. The site inherits the VLAN IDs specified for each TLS and a separate customer profile is configured for each TLS/site combination. Frames are transmitted across the correct TLS/VC LSP.





If a layer 2 site is linked to multiple TLSs, there must be no overlap between the VLAN IDs specified in the parent TLS definitions.

At minimum, a layer 2 site may contain a single port – represented by an Access interface on a Gateway (PE) device in the user interface. Service Activator configures the port as:

- An access port if the site receives untagged frames
- A trunk port if the site receives tagged frames

If Service Activator is to tag incoming frames at a layer 2 site, tagging is applied at the Access interface and access port configuration is applied. A VLAN definition is applied to the port making it part of the relevant VLAN.

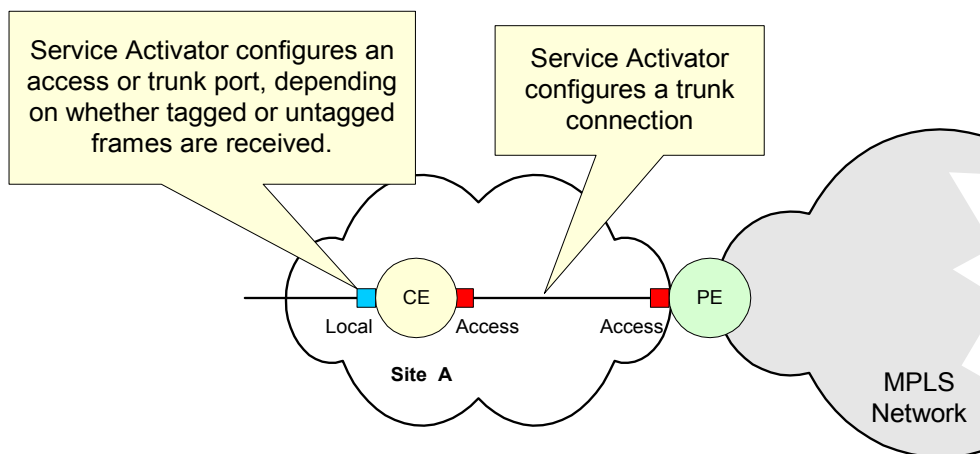
If the CE (Access) device is managed by the service provider, the site may also contain the CE device and at least one of its ports – represented by Access and Local interfaces in the user interface. Note the following:

- Only one Access interface on a CE device may be linked to a layer 2 site
- It is not essential to link the relevant Access interface to the layer 2 site in the user interface. If the CE device has an interface tagged with the Access role, Service Activator automatically uses that interface in its TLS configuration. Note, however, that if the device has two interfaces tagged with the Access role,

Service Activator cannot determine which interface to use and a fault will be generated.

- Any number of Local interfaces on a CE device may be linked to a layer 2 site
- Service Activator configures the Access interfaces on the CE and PE devices as trunk ports, creating a trunk connection. The CE device’s Local interfaces are configured as:
- Access ports if the site receives untagged frames
  - Trunk ports if the site receives tagged frames

If Service Activator is to tag incoming frames, tagging is applied at the CE device’s Local interface and trunk port configuration applied. A VLAN definition is applied to the port (Local interface) making it part of the relevant VLANs.



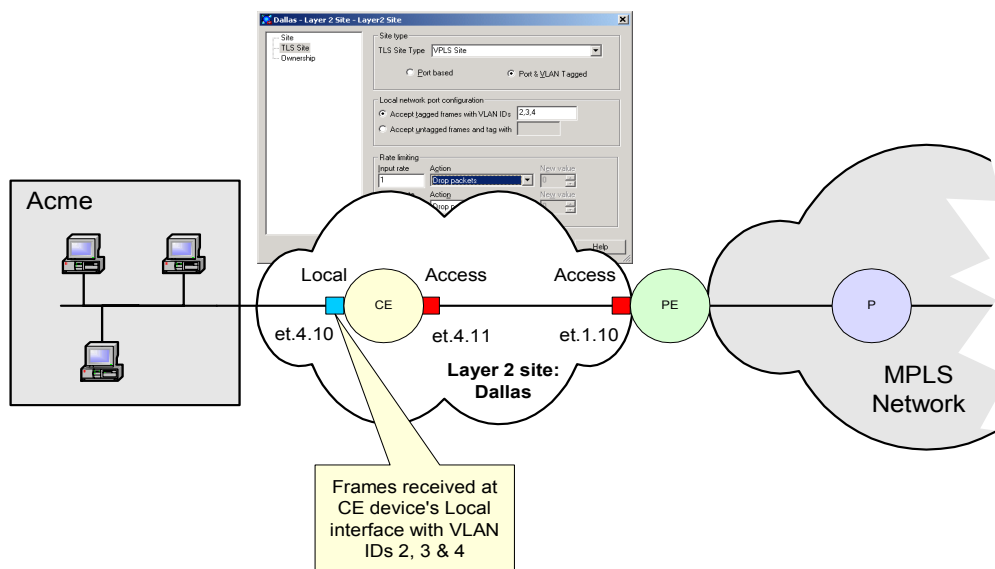
### Using a port in more than one port and VLAN-based TLS

A port may be linked to more than one port and VLAN-based layer 2 site, and so participate in multiple TLSs, provided:

- Tagging is not performed at the port
- The layer 2 sites the port is linked to specify a range of VLAN IDs
- There is no overlap of VLAN IDs between the layer 2 sites and the TLSs to which the port is linked

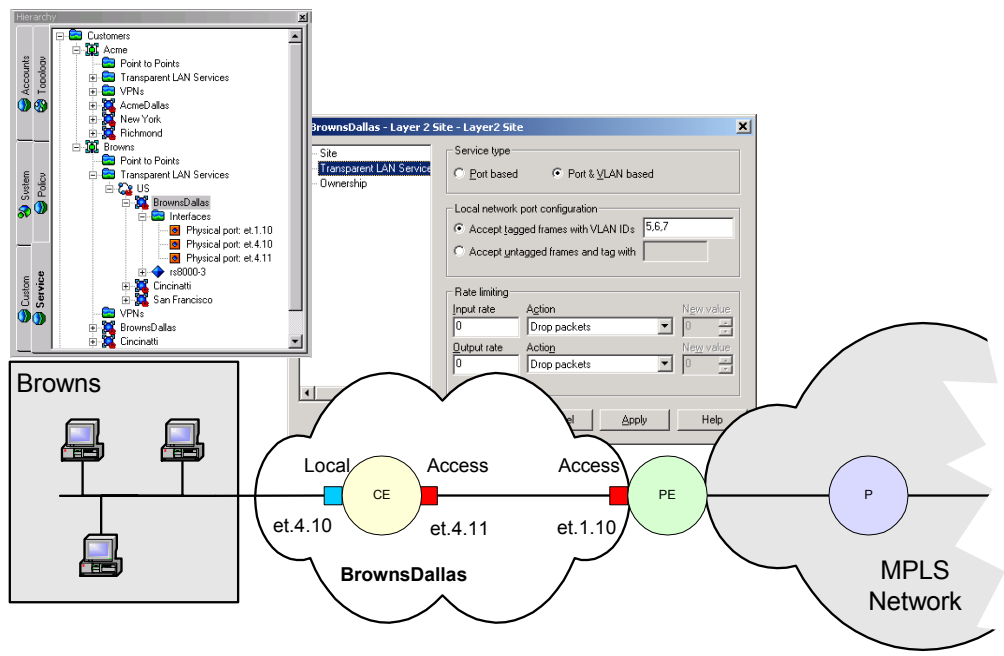
Note, however, that it is permissible to use the same VLAN IDs in different TLSs on the same PE device – that is, there may be overlap of VLAN IDs used on different ports provided they are part of different TLSs.

The following example shows a site setup for a port and VLAN-based TLS, where the service provider provides both the TLS and management of Ethernet VLAN IDs for a customer named Acme. The site contains both the Access interface on the PE device and Local and Access interfaces on the CE device. Incoming Ethernet frames are pre-tagged on reception by the CE device’s Local interface and are forwarded across the TLS.



The service provider wishes to use ports on the PE and CE device to provide an entirely separate TLS for another customer named Browns.

If there were no VLAN IDs specified in the Dallas layer 2 site definition – that is, the VLAN IDs were inherited from the TLS – the site could simply be linked to a TLS owned by a different customer. However, as VLAN IDs are specified at the site, the service provider must create a new site, specifying different VLAN IDs. The relevant ports can then be linked to the site to re-use the ports.



### Creating a TLS to support a Stacked VLAN

You can create a TLS object and Layer 2 sites to support stacked VLANs. In a stacked VLAN, packets arriving at the PE access port are already tagged with one or more VLAN IDs. Then an additional tag is added (a single VLAN ID) to identify them in the core VLAN. Upon exit, the stacked VLAN ID tag is removed, leaving the original customer-specified VLAN IDs. This preserves resources in the Core and allows the customer to use their own VLAN ID allocation scheme.

When created a stacked VLAN, select **Stacked VLAN, tagged with** on the **TLS Service** property page of the **Transparent LAN Service** dialog box. You can filter allowed customer VLAN IDs by specifying them in the adjacent field.

Set your Layer 2 sites up in a similar fashion, ensuring that your dialog box selections match the TLS.

For complete information on stacked VLANs and setting up VLANs in the core network using the **VLAN Activation Module**, refer to the **Service Activator on-line help**.

## Before setting up a TLS

### Manual pre-configuration

Some manual pre-configuration is required before setting up a TLS. For detailed information on manual pre-configuration, see the *Riverstone Device Driver Guide*.

#### PE devices

You must enable and start MPLS and LDP. In addition, an IGP routing protocol must be configured.

#### P devices

You must enable and start MPLS and LDP, and configure the IGP routing protocol.

### Discovering the network and assigning roles

When run, the discovery process finds all the P, PE and CE routers in the network. Service Activator stores these details in its database.

For information on running a discovery, see the *Network Discovery and Basic Setup* guide.

Note the following:

- All devices within the network must be correctly assigned system-defined roles, that is, PE routers must be classified as Gateway devices, P routers classified as Core devices and CE routers, if visible, as Access devices. The recommended way of assigning roles is by means of role assignment rules, which automatically assign roles during device discovery. If you do not use role assignment rules you need to assign a role manually for each device. For more information, see the *Network Discovery and Basic Setup* guide.
- All interfaces within the network must be correctly assigned system-defined roles:
  - On CE (Access) devices, the interface connected to the PE device must be classified as an Access interface. Interfaces connected to local segments must be classified as local interfaces.  
  
If a CE device is linked to a Layer 2 site, Service Activator automatically applies trunk port configuration to an interface tagged with the Access role (see [page 85](#))
  - On PE (Gateway) devices, the interface connected to the CE device must be classified as an Access interface. Interfaces connected to other PE devices or P (Core) devices must be classified as Core interfaces.

As for devices, we recommend you assign roles using role assignment rules. If you don't use role assignment rules you need to assign a role manually for each interface. For more information, see the *Network Discovery and Basic Setup* guide.

## Assigning devices to proxy agents

All devices that are to be managed by Service Activator must be assigned to a proxy agent. This is generally performed automatically during device discovery, but if devices are not assigned to the correct proxy agents you must assign them manually. For information on assigning devices to proxy agents, see the *Network Discovery and Basic Setup* guide.

## Setting devices to Managed

All devices to be configured by Service Activator need to have their status set to Managed. When devices are first discovered, their status is Unmanaged.

### To set all devices to Managed

- Select **Manage All Devices** from the network map's pop-up menu.

### To set an individual device to Managed

- Select **Manage Device** from the device's pop-up menu.

The device's color changes to reflect its new status. A managed device is represented by a green icon, an unmanaged device is represented by a blue icon.

## Setting up customers

A TLS is created for a customer – before you can create a TLS you must create the customer to whom the TLS belongs. For information on creating customers see [Setting up customers on page 20](#).

## Setting up the TLS

You create a TLS in Service Activator by creating a TLS object and associating layer 2 sites with the object. The sites that you link to a TLS mark the edge-points of the VLANs that are to be interconnected.

For information on planning a TLS, see [page 79](#).

## Creating a TLS

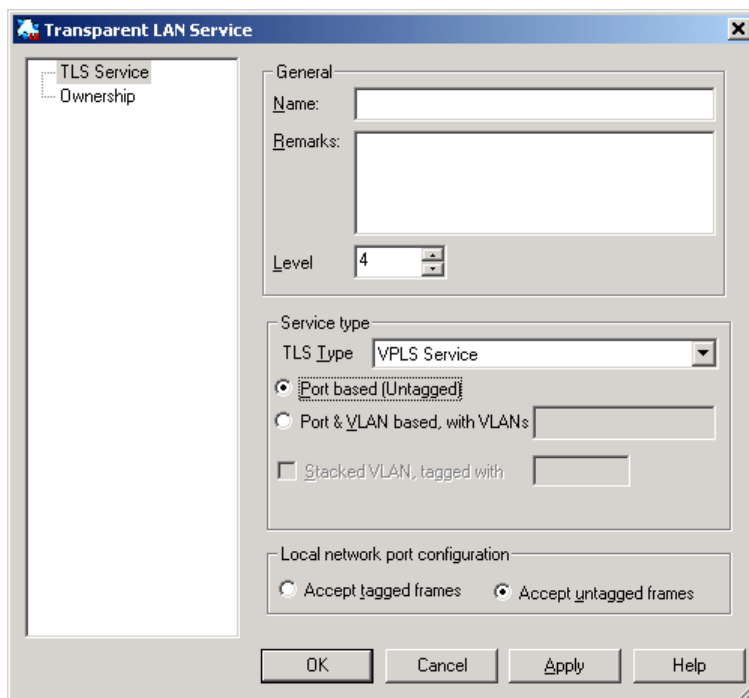
When creating a TLS object, you define broad characteristics of the TLS:

- For a port-based TLS, whether the TLS accepts untagged or tagged frames.  
Service Activator does not create VLANs or assign VLANs to any ports. It simply configures all ports in the TLS consistently as Trunk or Access ports.
- For a port and VLAN-based TLS, specify the VLAN ID or range of IDs that are to control access to the TLS.  
Service Activator creates VLANs on the devices and assigns ports to them. Each layer 2 site within the TLS may use all or a subset of VLAN IDs specified for the TLS. The created VLANs take their IDs from the layer 2 site to which the device’s ports are linked.

### To create a TLS

1. On the **Service** tab, open the relevant customer folder, select the **TLS** folder and select **Add Transparent LAN Service** from the pop-up menu.

The **Transparent LAN Service** dialog box opens.



2. On the **TLS Service** page:

- **Name:** specify a name for the TLS. The name may contain alphanumeric characters only, and may not include spaces or hyphens.
- **Remarks:** add any additional remarks (optional).
- **Level:** a level number for the TLS, in the range 0-7. This parameter is only used if a site is a member of more than one TLS and you are setting up QoS or access control on the TLS:
  - Rules are installed from all TLSs. Rules are installed in TLS level order, where rules from the TLS with the lowest level number are installed first and therefore evaluated first.
  - For PHB groups, up to a maximum of one PHB group is installed from the TLS with the lowest level number.

If two or more TLSs have the same priority level, a conflict is reported. By default the level is set to 4.

- **Service type:** specify the criteria on which access to the TLS is based:
  - **TLS Type:** select either **VLAN Service** or **VPLS Service** (select this for TLS on Riverstone)
  - **Port based:** forward traffic across the TLS according to incoming port number.
  - **Port & VLAN based, with VLANs:** forward traffic across the TLS according to incoming port number and the specified VLAN IDs.

VLAN IDs may be specified as a comma-separated list. The list may include ranges whose endpoints are separated by a hyphen. For example, 2,3,4-6. The allowable range of VLAN IDs is 2 to 4095.

When **Stacked VLAN, tagged with** is **unchecked**: specify the list of permitted VLAN IDs. Traffic will be passed on through the core with these IDs intact.

When **Stacked VLAN, tagged with** is **checked**: specify the list of permitted VLAN IDs. Packets with customer VLAN IDs matching the list are then tagged with the stacked VLAN IDs and passed across the core. On the other end, the stacked VLAN ID is stripped, leaving the original customer-specified VLAN IDs.

- **Stacked VLAN, tagged with:** Select this checkbox to select a Stacked VLAN. Use this option when the packets coming into the site from the CE are already tagged with a customer-specific VLAN ID. An additional tag is added as the packets move out of the site to the PE.

Provide the VLAN IDs that packets are to be tagged with.

**Note:** This checkbox is available only when **VLAN Service** is selected for the **TLS Type**.



- If the **Service type** is **Port based**, specify the **Local network port configuration**:
  - **Accept tagged frames**: transport only pre-tagged frames on the TLS (trunk port).
  - **Accept untagged frames**: transport only untagged frames on the TLS (access port).
- 3. If you wish to restrict access to the TLS object, select the **Ownership** page and specify the details – for information on setting ownership options, see the *Network Discovery and Basic Setup* guide.

## Setting up layer 2 sites

A layer 2 site defines an access point to a TLS:

- A port-based site may have a single port associated with it and be linked to a single port-based TLS
- A port and VLAN-based site may have ports on both the PE and the CE device associated with it

The site may be linked to multiple port and VLAN-based TLS objects provided there are no VLAN ranges specified at the layer 2 site and the site does not perform tagging.

For more information on port-based and port-based sites, see [page 80](#); for information on port and VLAN-based sites, see [page 82](#).

You can apply rate limiting to any site (see [Applying rate limiting to a layer 2 site on page 97](#)).

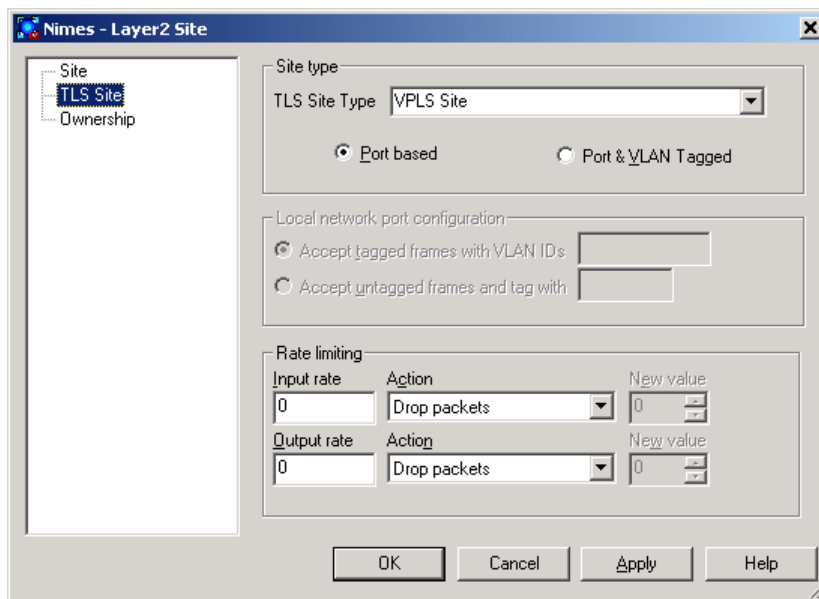
### To set up a layer 2 site

1. On the **Service** tab, right-click on the relevant customer folder and select **Add Layer2 Site** from the pop-up menu.

The Layer2 **Site** dialog box opens.

The screenshot shows a Windows-style dialog box titled "Layer2 Site". On the left side, there is a tree view with the following items: "Site", "TLS Site", and "Ownership". The main area of the dialog contains several input fields: "Name:" (a single-line text box), "Remarks:" (a multi-line text area), "Account Ref:" (a single-line text box), "Contact:" (a single-line text box), "Address:" (a multi-line text area), "Tel:" (a single-line text box), "Fax:" (a single-line text box), and "E-mail:" (a single-line text box). At the bottom of the dialog, there are four buttons: "OK", "Cancel", "Apply", and "Help".

2. On the **Site** page, specify an identifying **Name** for the site, and any additional comments. You can set up account and contact information if required, but this is optional.
3. On the **Transparent LAN Service** page, select the **Service type**:
  - **TLS Site Type**: select either **VLAN Service** or **VPLS Service** (select this for TLS on Riverstone)
  - **Port based**: create a port-based TLS on the port.
  - **Port & VLAN Tagged**: create a port and VLAN-based TLS on the port.



4. If the **Service type** is **Port & VLAN Tagged**, use the **Local network port configuration** frame to specify how the port handles incoming frames:
  - **Accept tagged frames with VLAN IDs**: accept frames tagged with the specified VLAN IDs (trunk port configuration).

VLAN IDs may be specified as a comma-separated list. The list may include ranges whose endpoints are separated by a hyphen. For example, 2,3,4-6. Note that the specified ID must also be within the range of VLANs specified on the TLS.

If no VLAN IDs are specified for the site, Service Activator uses those specified for the TLS.

Incoming frames belonging to VLANs that are outside the specified range are dropped.

- **Accept untagged frames and tag with:** accept untagged frames and tag them with the specified VLAN ID (access port configuration).

Note that the specified ID must also be part of the range of VLANs specified on the TLS.

VLAN ID 1 is reserved for the default VLAN – the VLAN that all unconfigured ports belong to by default. Any traffic tagged with VLAN 1 can be transported to another port in VLAN 1 only.

5. If required, specify rate limiting parameters for the site (see [Applying rate limiting to a layer 2 site on page 97](#)).
6. If you wish to restrict users' access to the layer 2 site object, select the **Ownership** page and specify the details – for information on setting ownership options, see the *Network Discovery and Basic Setup* guide.

## Associating a physical component with a layer 2 site

Every layer 2 site must have, at minimum, the port on the relevant PE device linked to it. If you are creating a port and VLAN-based site and the service provider has control of the CE device, this may also be linked to the site.

For more information, see [Planning the TLS on page 79](#).

### To link a PE access interface to a layer 2 site

- Drag and drop the appropriate access interface on the PE (gateway) device on to the site.

## Linking a CE router to a layer 2 site

If Service Activator has visibility of site CE routers, you should also link the CE device to the site and one Access and one or more Local interfaces. This is only required if the service provider is offering a fully-managed TLS and has complete visibility of the customer's devices.

### To link the CE device

- Drag and drop the CE device on to the site and at least one access and one local interface.

## Linking sites to a TLS

You define the access points to the TLS by linking the appropriate customer sites to the TLS object:

- A port-based site may be linked to a port-based TLS

- A port and VLAN-based site may be linked to a port and VLAN-based TLS

If a port and VLAN-based layer 2 site does not define any VLAN IDs, but instead inherits them from the TLS to which it is linked, the site may be linked to another TLS. However, the range of VLAN IDs on the various TLSs the site is linked to must not overlap.

#### To link a site to a TLS

- Drag and drop the layer 2 site object onto the TLS to create a link.

## Applying rate limiting to a layer 2 site

Rate limiting constrains outbound traffic to a particular bandwidth and is commonly used to control access to the core network. It can be used to regulate the flow of traffic in order to avoid the congestion that can occur when transmitted traffic exceeds the access speed of the remote interface.

You can apply rate limiting to one or more of the sites associated with a TLS, specifying the maximum input and output speed. Service Activator applies the specified parameters as follows:

- If the site contains only the PE access port, rate limiting is applied to this port
- If the site contains both the PE access port and the CE device and relevant interfaces, rate limiting is applied to the CE device's local port

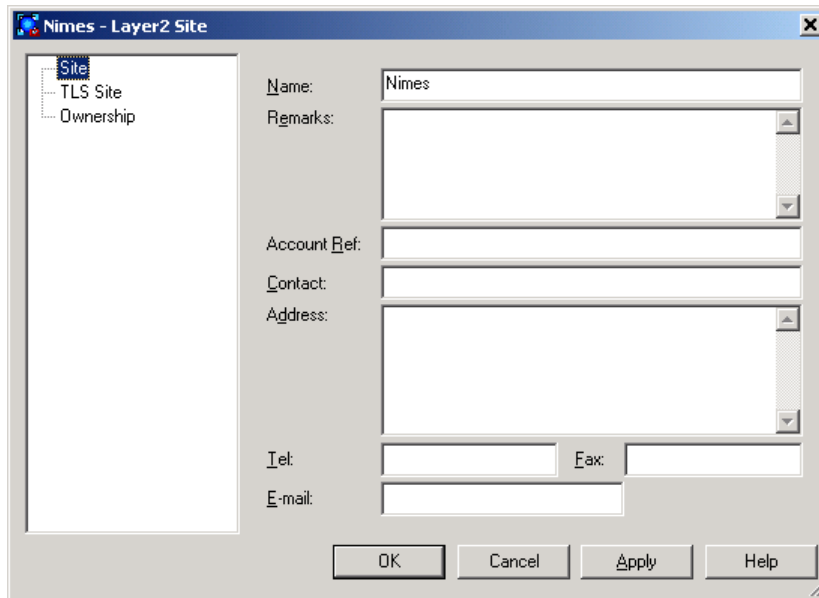
The maximum possible input and output speed for rate limiting is interface-dependent. For information on checking an interface's rate limiting range, consult the Riverstone documentation.

Note that in order to configure port rate limiting policies for input ports, you must first enable the aggregate rate limiting mode on the line card. For more information, see the *Riverstone Device Driver Guide*.

#### To apply rate limiting to a layer 2 site

1. On the **Service** tab, open the relevant customer folder, open the **Transparent LAN Services** folder and the relevant TLS and open the relevant site.

The Layer2 **Site** dialog box opens.



2. On the **Transparent LAN Service** page specify rate limit values for incoming and outgoing traffic on the site's port:

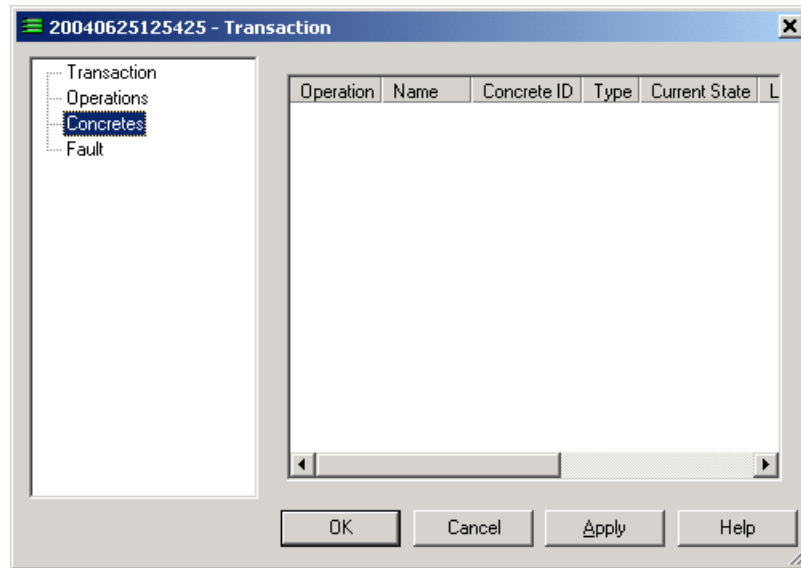
Option	Description
Input rate	The rate limit value for incoming traffic in bits per second.
Action	The action to take when the rate limit is exceeded: <ul style="list-style-type: none"> <li>■ <b>Drop packets:</b> packets are discarded</li> <li>■ <b>Reduce priority:</b> lower the packet's priority</li> <li>■ <b>Rewrite TOS:</b> rewrite the ToS byte to the value specified in the <b>New value</b> field</li> <li>■ <b>Rewrite TOS precedence:</b> rewrite the packet's ToS precedence to the value specified in the <b>New value</b> field</li> </ul>

Option	Description
New value	<p>The TOS bit or TOS precedence value to apply if the input rate limit is exceeded.</p> <p>This field is only enabled if Rewrite TOS or Rewrite TOS precedence is selected in the Action field:</p> <ul style="list-style-type: none"> <li>■ TOS bit may be in the range 0-255</li> <li>■ TOS precedence may be in the range 0-7</li> </ul>
Output rate	The rate limit value for outgoing traffic in bits per second.
Action	<p>The action to take when the rate limit is exceeded:</p> <ul style="list-style-type: none"> <li>■ Drop packets</li> <li>■ Reduce priority</li> <li>■ Rewrite TOS</li> <li>■ Rewrite TOS precedence</li> </ul>
New value	<p>The TOS bit or TOS precedence value to apply if the input rate limit is exceeded.</p> <p>This field is only enabled if Rewrite TOS or Rewrite TOS precedence is selected in the Action field:</p> <ul style="list-style-type: none"> <li>■ TOS bit may be in the range 0-255</li> <li>■ TOS precedence may be in the range 0-7</li> </ul>

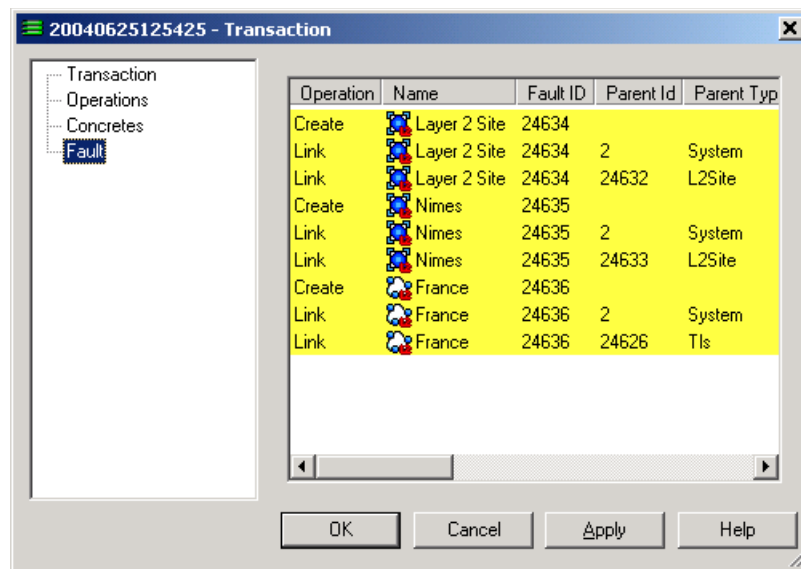
## Implementing the TLS

Once the site and TLS details are set up, the entire configuration can be applied by committing the transaction.

When you commit the transaction, any concrete TLSs that will be created are listed in the **Concretes** page of the Transaction dialog box.



Any validation errors are reported in the Transaction dialog box and the **Current Faults** pane.



If you wish to cancel the transaction after reviewing the concrete TLSs that will be created and the faults generated by the transaction, click **Cancel**.



If you wish to proceed with the transaction, click **OK**. Configuration details are sent to the proxy agent and on to the appropriate device drivers. For information on committing a transaction, see *Network Discovery and Basic Setup*.

After committing the transaction, you can check the configuration that has been applied to the routers by checking the device logs and by using Telnet. For more information, see the *Administrator's Guide*.

## Viewing implemented TLSs

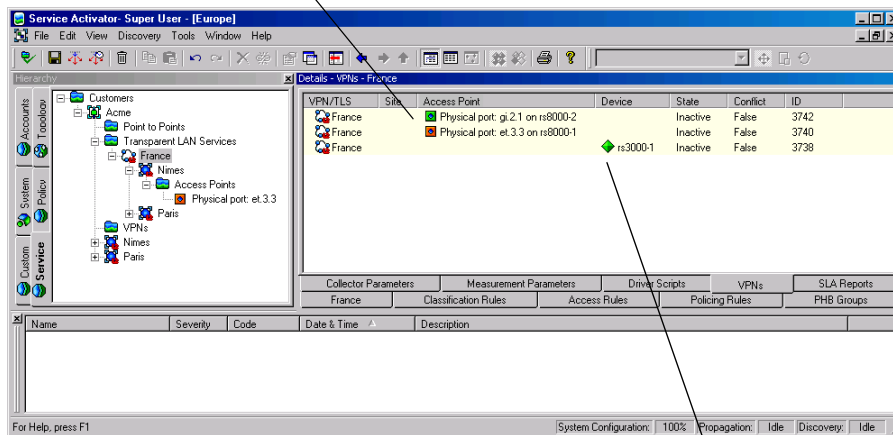
You can view a list of the TLSs that have been propagated to the network and installed on an interface.

On the CE device, a concrete TLS represents the VLANs that have been created and assigned to the device's ports. Note that you cannot view the concrete VLANs that apply to a specific port on a CE device.

On the PE device, a concrete TLS represents the application of a customer profile to a port.

By viewing concrete TLS details for a TLS object, you can view the points in the network at which TLS or VLAN configuration has been applied.

On a PE device, a concrete TLS represents the application of a customer profile to a port



On a CE device, a concrete TLS represents the VLANs that have been created and assigned to the device's ports

### To view implemented TLS details

- Double-click on the relevant object from the hierarchy tree or the topology map:  
You can view concrete TLSs for:
  - A CE device – represents VLAN configuration that applies to all ports
  - A port on a PE device – represents the application of a customer profile to a specific port
  - A TLS – represents the points in the network at which TLS or VLAN configuration has been applied
- In the details pane, select the **VPNs** tab to view details of the TLS configuration that applies to a port on a PE device, or a CE device.

## Chapter 4

# Setting Up Point-To-Point Connections

This chapter describes how to configure point-to-point connections on Juniper M-series devices. The chapter:

- Provides an overview of point-to-point connections and Circuit Cross Connects (CCCs)
- Describes the manual pre-configuration tasks associated with CCCs
- Explains how to create and delete CCCs
- Describes the QoS elements that can be applied to CCCs

## Overview

Circuit Cross Connect (CCC) is a proprietary Juniper feature that allows you to configure transparent connections between two circuits. A circuit can be any of the following:

- Frame Relay DLCI
- ATM VC
- Point to Point Protocol (PPP) interface
- Cisco High-level Data Link Control (HDLC) interface
- Virtual Local Area Network (VLAN)

Using CCC, packets from one circuit are forwarded to another with only the level 2 address changed, at most.

Juniper M-series devices support three types of CCCs:

- Layer 2 switching – essentially provide Layer 2 switching between two logical interfaces of the same type on the same device
- MPLS tunneling – connects two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the channel
- LSP stitching – stitches together two LSPs

Service Activator's Juniper M-series device driver supports Layer 2 switching and MPLS tunneling CCCs.

For all CCC connections that connect interfaces, the interfaces must be of the same type; that is, ATM to ATM, Frame Relay to Frame Relay, PPP to PPP, or Cisco HDLC to Cisco HDLC.

If the interfaces are on different devices, the connection is created by two MPLS Label Switched Paths (LSPs) across the network – one for each direction (cross-connections are duplex, but LSPs are simplex).

If the interfaces are on the same device, the connection is achieved by layer 2 switching across the device's backplane.

If a Layer 3 network configuration has been applied it is unsuitable to create point-to-point circuit cross connections.

If a Layer 2 circuit (such as CCC encapsulation) is configured on a logical interface, it is unsuitable to configure another protocol family on the same interface.

In order to configure any type of CCC, MPLS must be enabled on all the routers along the path of the CCC.

For MPLS tunneling CCCs, it is also recommended that you configure MPLS and RSVP on all devices that may potentially be used in the LSP.


Refer to the *Juniper Device Driver Technical Note* for further information about configuring encapsulation, CCC connections, and MPLS circuits on Juniper M-series devices.

## Layer 2 switching CCC

A layer 2 switching CCC enables you to join two circuits, effectively configuring the interdomain router as a switch.

For example, in the following diagram, a CCC connects two Frame Relay circuits. Router B acts as the Frame Relay switch, transparently switching packets between Router A and Router C. The only processing performed by Router B is to translate DLCI 600 to 750.

**Figure 2: Layer 2 switching CCC**




## MPLS tunneling CCC

MPLS tunneling CCCs connect two distant interface circuits of the same type. This creates an MPLS tunnel that uses an LSP as the conduit.

For example, in the following diagram, a CCC connects two ATM access networks through an IP backbone. CCC establishes an LSP tunnel between the two domains. ATM traffic from one network is tunneled across a SONET backbone to the second network using an MPLS LSP.

**Figure 3: MPLS tunneling CCC**



## Manual pre-configuration

Some manual pre-configuration is required for MPLS tunneling CCCs.

- A loopback interface on the LSP's ingress and egress devices must be configured to provide termination points for the CCC.
- Service Activator automatically configures MPLS on the LSP's ingress and egress routers. RSVP must be manually configured on these routers.
- CCC core routers and PE core interfaces must be running MPLS, RSVP, and an IGP.

The LSP is dynamic and packet routing may change depending on traffic congestion. We therefore recommend that you configure MPLS and RSVP on all devices that may potentially be used in the LSP.

For configuration details, see the *Juniper Device Driver Technical Note* and the *MLPS LSP Module* online help.

### Manual pre-configuration: CCCs on Ethernet interfaces

If any physical interface encapsulation incompatibilities pre-exist on the router, Service Activator detects them when the device driver is building a new configuration for CCCs. An error is displayed in the UI, and you are given the option to manually correct the interface encapsulation.

In order to expedite the configuration process, ensure that the following manual configuration exists on the router:

1. For 802.1Q VLANs and/or VLAN-based circuits, ensure that vlan-tagging is enabled on physical interfaces, and that each logical subinterface has a VLAN ID configured.
2. For physical Ethernet interfaces to be used in port-based CCCs, ensure that there is no vlan-tagging and either only unit 0 or none of logical subinterfaces are present.

## Provision sub-interfaces for CCCs

You must first create the sub-interfaces and CCC endpoints before linking them to the CCC object. The procedures are the same as for creating sub-interfaces and endpoints for Layer 2 Martini connections.

1. To create the sub-interfaces for new CCCs, follow the procedure [Provisioning endpoints \(VC IDs\) for a Martini L2 connection on page 70](#).
2. To provision the endpoints for new CCCs, follow the procedure [Provisioning endpoints \(VC IDs\) for a Martini L2 connection on page 70](#).

### Example values for each endpoint of an ATM CCC:

- **Sub-Interface:** the ID of the sub-interface where the CCC will exist.
- **VPI / VCI:** the virtual path identifier and the virtual channel identifier.
- **Max VCs:** the maximum number of virtual channels that the interface can support.

**Example values for each endpoint of a Frame Relay CCC:**

- **Sub-Interface:** the ID of the sub-interface where the CCC will exist.
- **DLCI:** the data link connection identifier. The value must be in the range 512 to 1022.

**Example values for each endpoint of an Ethernet VLAN CCC**

- **Sub-Interface:** the ID of the sub-interface where the CCC will exist.
- **VLAN Id:** the virtual local area network number. The value must be between 512 and 1023.

## Creating a CCC

CCCs are created on a per customer basis. For information on creating customers see [Setting up customers on page 20](#).

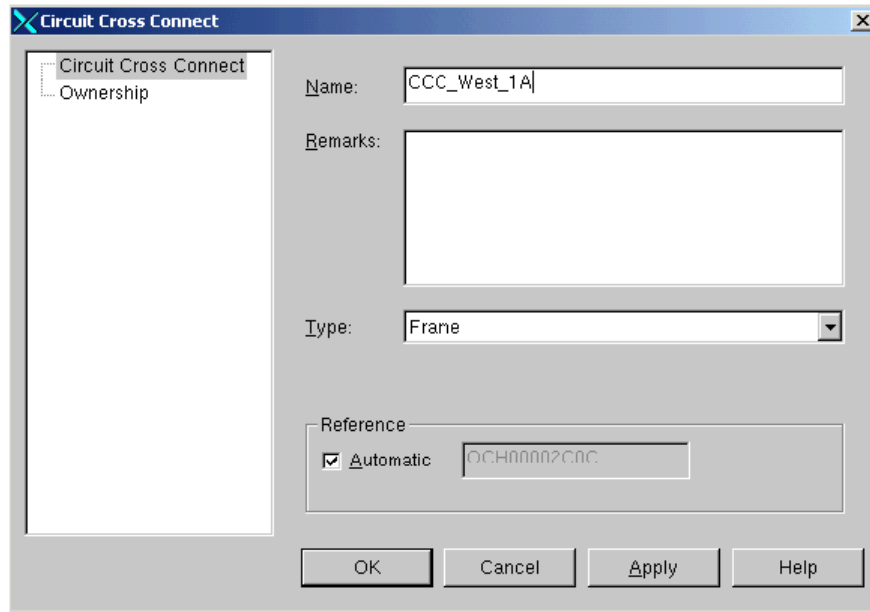
When connecting ATM, Frame Relay or VLAN circuits with a CCC you must define additional parameters to be used when encapsulating packets at the CCC endpoints.

**To create a CCC object**

1. Select the **Service** tab and open the **Customers** folder from the hierarchy pane.
2. Under the required customer, right click on the **Point to Points** folder, then select **Add Circuit Cross Connect** from the pop-up menu.



The **Circuit Cross Connect** dialog box opens.



3. Enter the following details.
  - **Name:** name of the CCC
  - **Remarks:** optional remarks
  - **Type:** the type of CCC, the options are:
    - **ATM AAL5**
    - **ATM Cell**
    - **Ethernet**
    - **Ethernet VLAN**
    - **Frame**
    - **HDLC**
    - **PPP**
  - **Reference:** The base name of the LSPs. This setting is only relevant to MPLS tunneling CCCs between remote interfaces:
    - **Automatic:** If selected, Service Activator automatically generates LSP names. Clear the checkbox if you wish to define the LSPs name.
4. Click **Apply**.

Service Activator creates a new CCC object and lists it in the **Point to Points** folder.

## Associating interfaces with a CCC

### Concepts around Interfaces and Endpoints

To create the CCC connection, you associate CCC endpoints with a CCC by dragging the endpoints and dropping them on the CCC object. The definition of the CCC endpoints depends on the type of CCC. For port-based CCCs (Ethernet and Frame Relay), the CCC endpoint is the physical interface. For all other types of CCCs, the CCC endpoints are the Layer 2 circuit IDs (as defined by frame relay DLCI IDs, Ethernet VLAN IDs, or ATM VPI/VCI IDs) configured under the logical interface.

If a logical interface and virtual circuit ID are already configured on an interface, you can specify this circuit ID as the CCC's endpoint.

(In this context, a logical interface is the sub-interface of a physical interface. Virtual circuit IDs such as DLCIs, VLANs, or VPI/VCI are defined under these logical interfaces.)

If logical interfaces do not exist, you can create a new logical interface and Layer 2 circuit using the Add Provisioned Sub-interface feature in the user interface. The new objects are displayed in the user interface on device rediscovery. Device rediscovery updates the state value and enables the new objects to be used in a CCC or Martini connection.

If the interfaces linked to the CCC object are on the same device, Service Activator configures a Layer 2 Switching CCC. If the interfaces are on different devices, Service Activator configures an MPLS Tunneling CCC (one CCC and two LSPs).

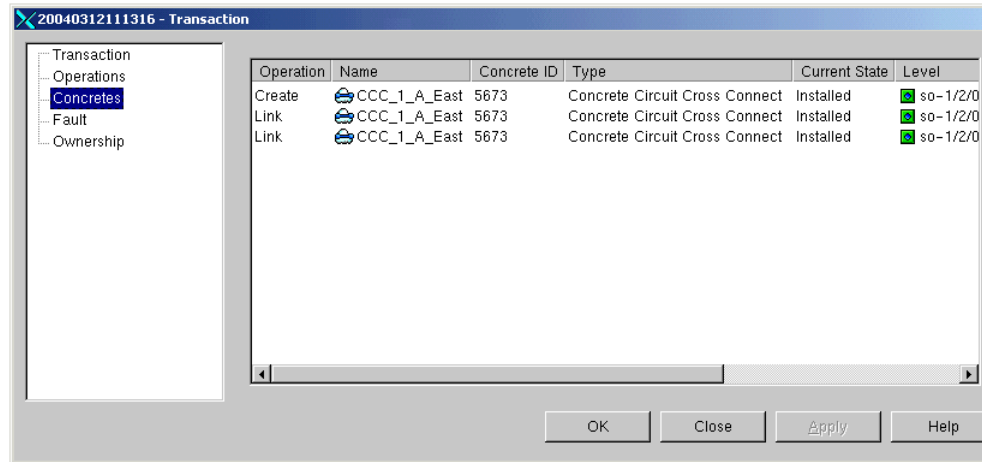
### To associate interfaces with a CCC

- Ensure that role **Access** is assigned to CCC endpoints, as well as their parent objects (logical and physical interfaces).
- Drag each endpoint and drop on the appropriate CCC object. (Alternatively, use the Copy and Paste Link commands.)

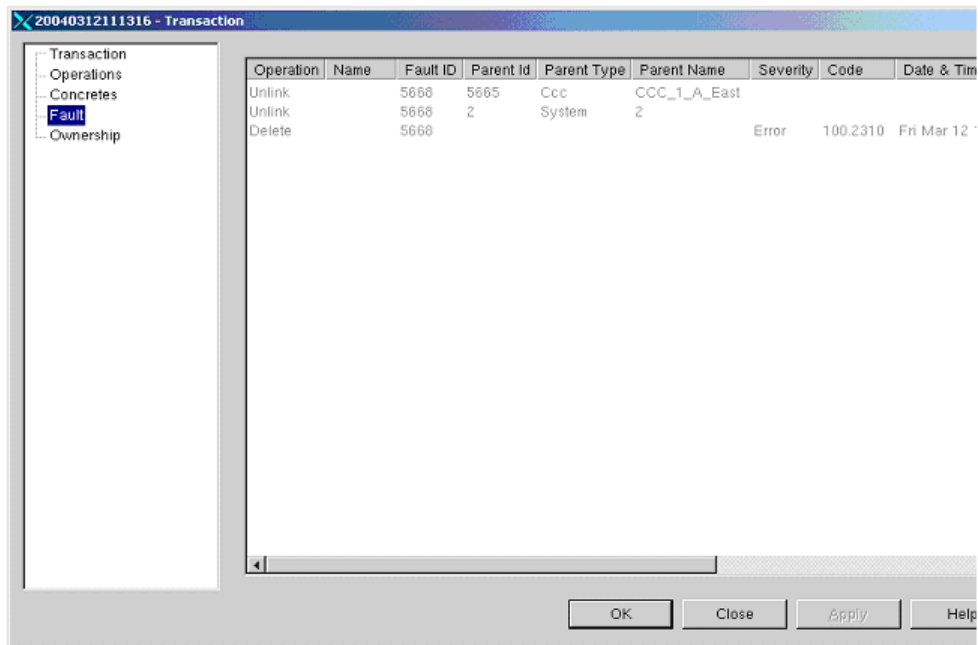
## Implementing CCCs

Once the CCC details are set up and the relevant devices are managed, the entire configuration can be applied by committing the transaction.

When you commit the transaction, any concrete CCCs that will be created are listed in the **Concretes** page of the Transaction dialog box.



Any validation errors are reported in the **Fault** page of the Transaction dialog box and the **Current Faults** pane.



If you wish to cancel the transaction after reviewing the concrete CCCs that will be created and the faults generated by the transaction, click **Cancel**.

If you wish to proceed with the transaction, click **OK**. Configuration details are sent to the proxy agent and on to the appropriate device drivers. For information on committing a transaction, see *Network Discovery and Basic Setup*.

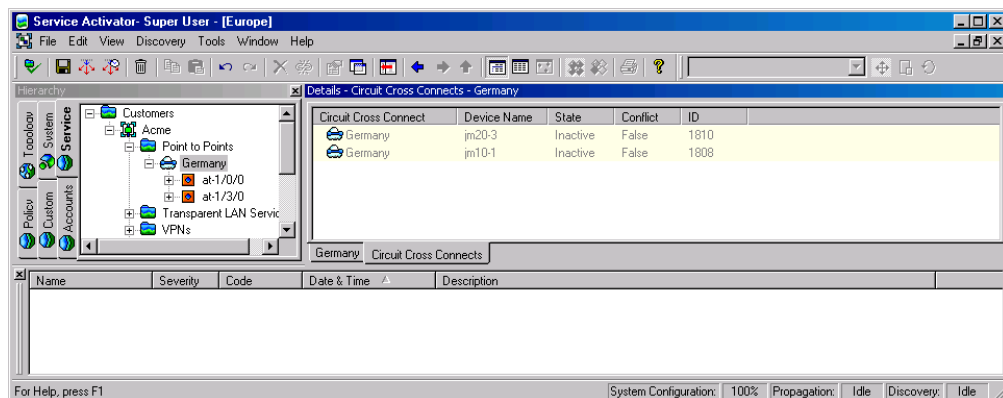
After committing the transaction, you can check the configuration that has been applied to the routers by checking the device logs and by using Telnet. For more information, see the *Administrator's Guide*

## Viewing implemented CCCs

You can view a list of the CCCs that have been propagated to the network and installed on devices.

### To view implemented CCC details

1. Double-click on the CCC from the hierarchy tree or the topology map.
2. In the details pane, click on the **Circuit Cross Connects** tab to view the CCC's implementation details. All concrete CCCs appear on a yellow background.



## Deleting CCCs

A CCC is deleted when one of its two endpoints is unlinked or deleted. The other circuit endpoint can be preserved for relinking, so that setting up the next CCC is faster. Alternatively, you can delete both endpoints to remove the CCC.

### To delete a CCC

1. Select the **Service** tab and open the **Customers** folder from the hierarchy pane.

2. Under the required customer, open the **Point to Points** folder, then select the folder of the CCC you want to delete.
3. Click the delete button in the main toolbar.

The CCC is deleted only when the transaction that holds the CCC deletion has been propagated to the network.

## Applying QoS to CCCs

If you wish to apply QoS to a CCC you must apply policy to the interfaces that are associated with the CCC. You cannot associate rules and PHB groups with a CCC object.

### To apply QoS to a CCC

1. Select the **Point to Points** folder.
2. Double-click the CCC you require and select the interface to which you want to apply policy.
3. From the interface's pop-up menu, select **Add Access Rule** or **Add PHB Group**.

For additional information, refer to the *Juniper Device Driver Technical Note*.

## PHB groups

You can apply the following mechanisms to interfaces within a CCC:

- **WRR** – WRR is a mechanism for allocating bandwidth to queues so that higher priority applications have more bandwidth than lower priority applications when the network is congested. You can set the weight to be allocated to each CoS. For information on configuring WRR on Juniper devices, see the *Juniper Device Driver Technical Note*.

The **juniper.policy** file is supplied with Service Activator to support WRR configuration. For further information see the *Juniper Device Driver Technical Note*.

- **Rate Limiting** – Rate limiting or traffic shaping, constrains specific outbound traffic to a particular bandwidth. It is commonly used to control access to the core network. It can be used to regulate the flow of traffic in order to avoid the congestion that can occur when transmitted traffic exceeds the access speed of the remote interface.

For each CoS included in a PHB group, you can set an average transmission rate, a burst rate and a burst interval. For further information on implementing rate limiting see the *Juniper Device Driver Technical Note*.

For information on configuring PHB groups, see *Configuring Policy Services*.

## Access rules

Access rules (or filters) are used to provide network security. Identified traffic can be explicitly denied or permitted access to the network.

Traffic can be identified by a combination of source and destination IP address or account and traffic type – for example, source or destination port. The traffic affected by the rule can be defined within the rule itself or based on one or more pre-defined traffic classifications.

Depending on how your system is configured, you may need to create one or a number of access rules for each traffic type to be managed. By grouping traffic types into traffic classification groups, you may be able to minimize the number of rules required. For information on defining traffic classifications and classification groups, see *Configuring Policy Services*

Access rules can only be used on devices with Internet Processor II ASIC. For more detailed information see the *Juniper Device Driver Technical Note*

## Chapter 5

# IPsec VPNs

This chapter explains how to configure point-to-point IPsec (Internet Protocol Security) and IP/GRE tunnels. It also describes concepts about implementing VRF-Aware IPsec. This chapter includes:

- Provisioning of IP/GRE tunnels with support for OSPF, EIGRP and static routing.
- Provisioning of IPsec tunnels
  - Support for IKE Pre-shared keys
  - Support for DES, Triple-DES, and AES Encryption
  - Support for MD5 and SHA-1 Packet Authentication
  - Support for Encapsulating Security Payload (ESP)
- Provisioning of IPsec over IP tunnels
- Removal of provisioned IP/GRE tunnels and IPsec Tunnels
- Configuring of Default IPsec and IP/GRE Tunnel Options that can be associated with a specific customer or become global default options

For instructions on how to install the IPsec VPN module, refer to *Service Activator Setup Guide*.

## First Time User Setup

### Screen Resolution

For best results, it is recommended that you run the IPsec application at a screen resolution of 1024 by 768 pixels.

## Oracle Database Login

The IPsec application requires login information to communicate with the Oracle database. The information includes the Oracle user name and password, the IP

Address of the database and the database SID. This information is normally gathered at install time, but you may see the dialog box shown below to get these login parameters.



If the dialog box appears, perform these steps to save the options:

1. Enter values in all the fields.
2. Click **OK** to continue to the next screen of the application or click **Cancel** to exit the application.

Perform this procedure once per machine. The system stores this information in the properties file `osa.properties`, which is read when needed. The password field is encrypted in the file for security reasons. This file is located in `C:\Program Files\Oracle Communications\IP Service Activator\modules\config`.


## Service Activator Login

The IPsec application requires user login information to communicate with the Service Activator object model. The dialog box prompts for your login name, login password, the IP Address of the machine running the Component Manager, and the port number. If the Login Parameters dialog box appears, perform these steps to save this information:

1. Enter values in all the fields.
2. Click **OK** to continue to the next screen of the application or click **Cancel** to exit the application.



By default, the IP Address for the Oracle database is pre-filled along with the port number; these can be changed if required. This information is stored in the database, it only has to be entered once per platform login name i.e. for Windows, the userid used to log in to Windows. This permits multiple users to access the IPsec application and to access it on different machines.



**Note:** To access the IPsec VPN functionality in the Service Activator GUI, ensure that the **Allow concurrent logins** option is enabled for the Service Activator user. For more information on how to do this, refer to the *Service Activator Administrator's Guide*.

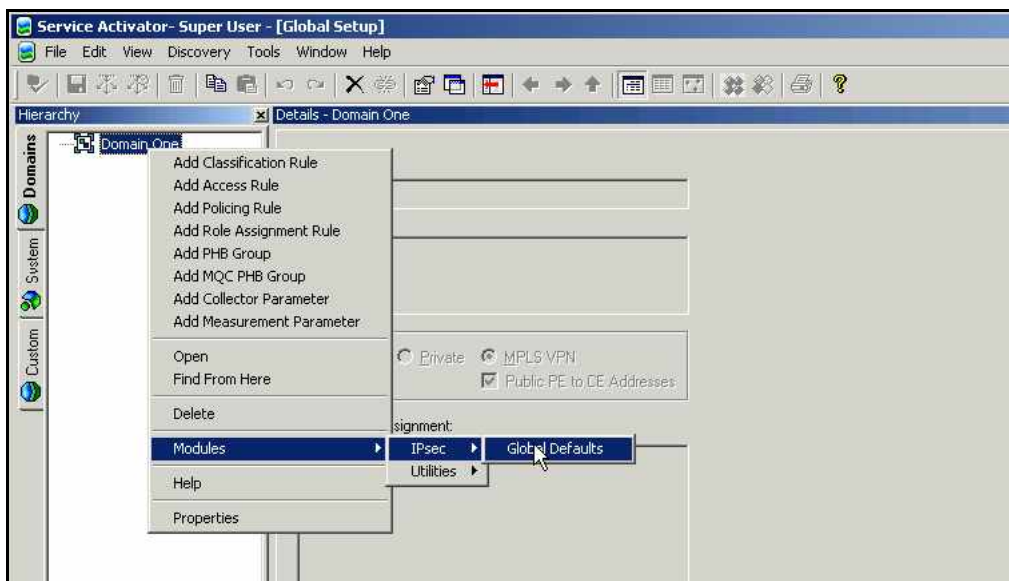
## Creating Global Default Options

The purpose of global default options is to reduce the configuration effort each time a user creates an IPsec or IP/GRE tunnel. There is only one set of Global default options for the entire application. Typically only one user should control what the global defaults are. Customer default options take precedence over Global default options, so Global default options will not be used if Customer default options exist for a customer that you are provisioning.

Default options are used for defaulting GUI fields when creating IPsec VPNs and IP Tunnels. If default settings are changed, it will have no effect on previously provisioned IP Tunnels and IPsec VPNs.

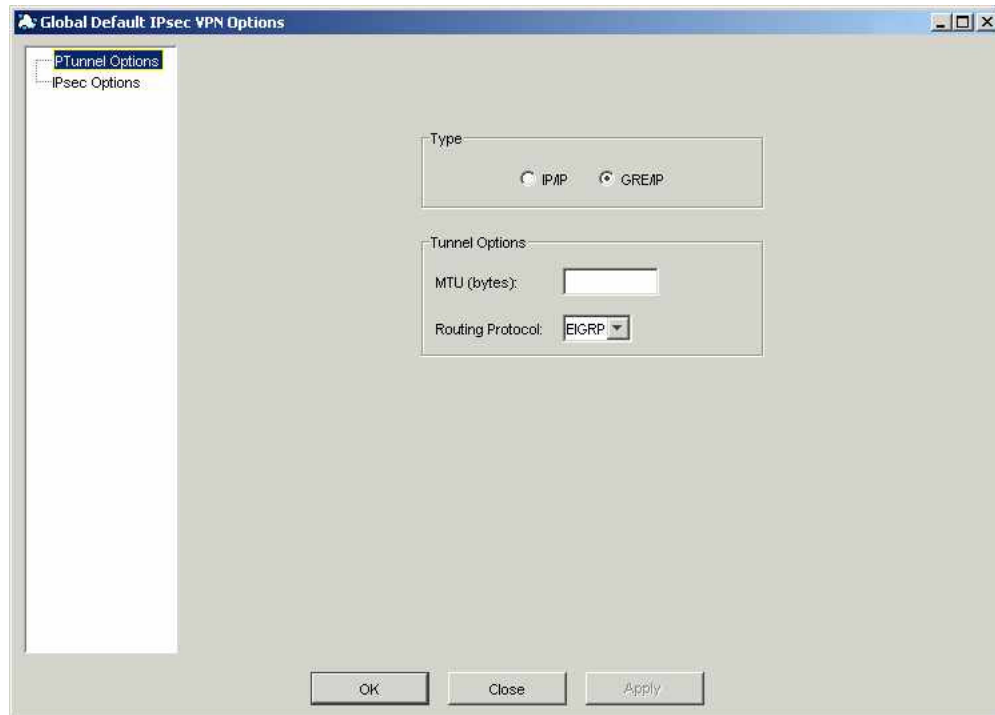
Perform this procedure to create global default options for IP/GRE Tunnels and IPsec VPN Tunnels.

1. Open the Service Activator GUI and right-click on any domain object. From the drop-down, select **Modules**, then **IPsec**, then **Global Defaults**.

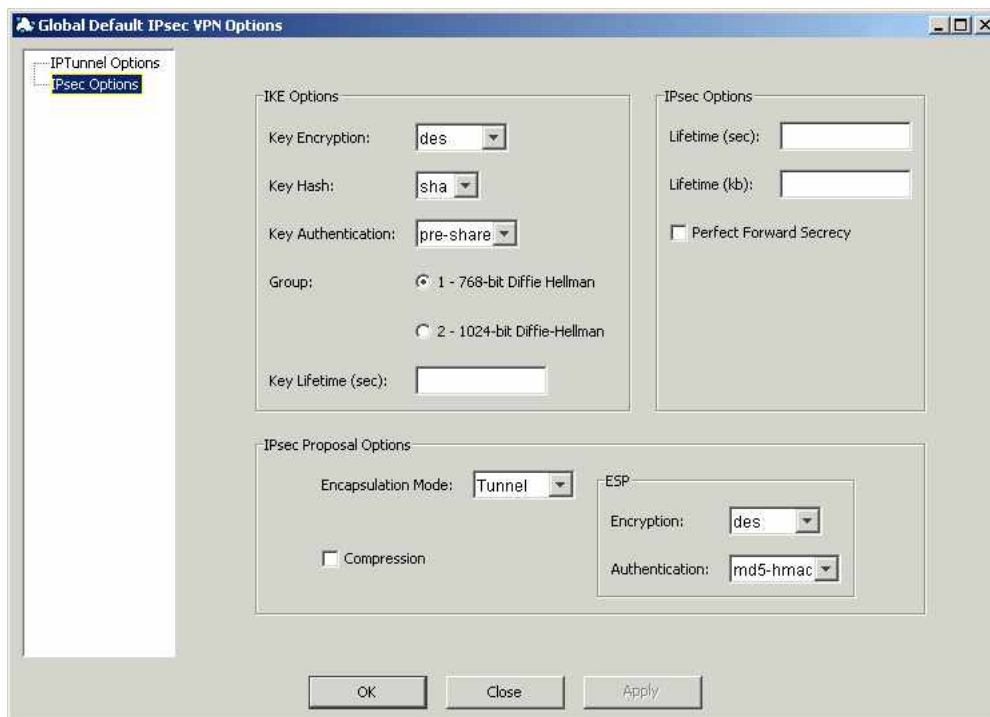


2. If this is your first time using the IPsec application, you are prompted for login information for Service Activator. Please see [Service Activator Login](#) on page 116 for more information.

The Global Default IPsec VPN Options dialog box appears.



3. Enter the IP Tunnel default options you wish to make as global defaults.
4. Click on IPsec Options to display its property page.



5. Enter the IPsec default options you wish to make as global defaults, then click 'OK'.

The button actions are as follows:

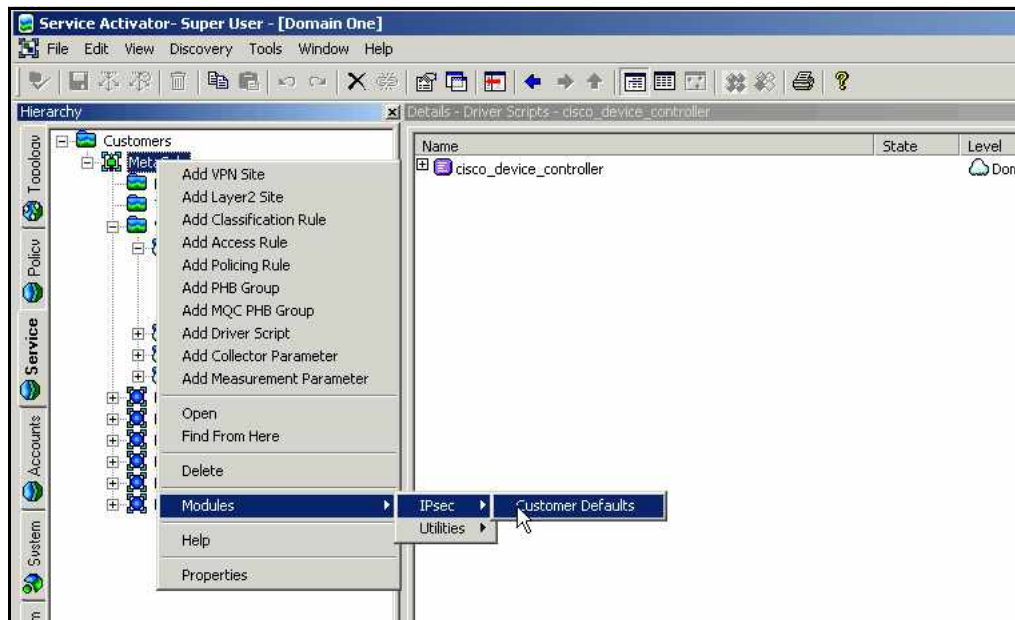
- 'Close' – exits
- 'Cancel' – exits without saving changes
- 'Apply' – saves any changes
- 'OK' – saves any changes and exits

## Creating Customer Default Options

Customer default options application allows users to assign default values to certain settings so that less configuration is needed each time a user creates an IPsec or IP/GRE tunnel. Customer default options take precedence over Global default options, so Global default options are used if Customer default options do not exist. This procedure outlines how to create customer default options for IP/GRE Tunnels and IPsec VPN Tunnels.

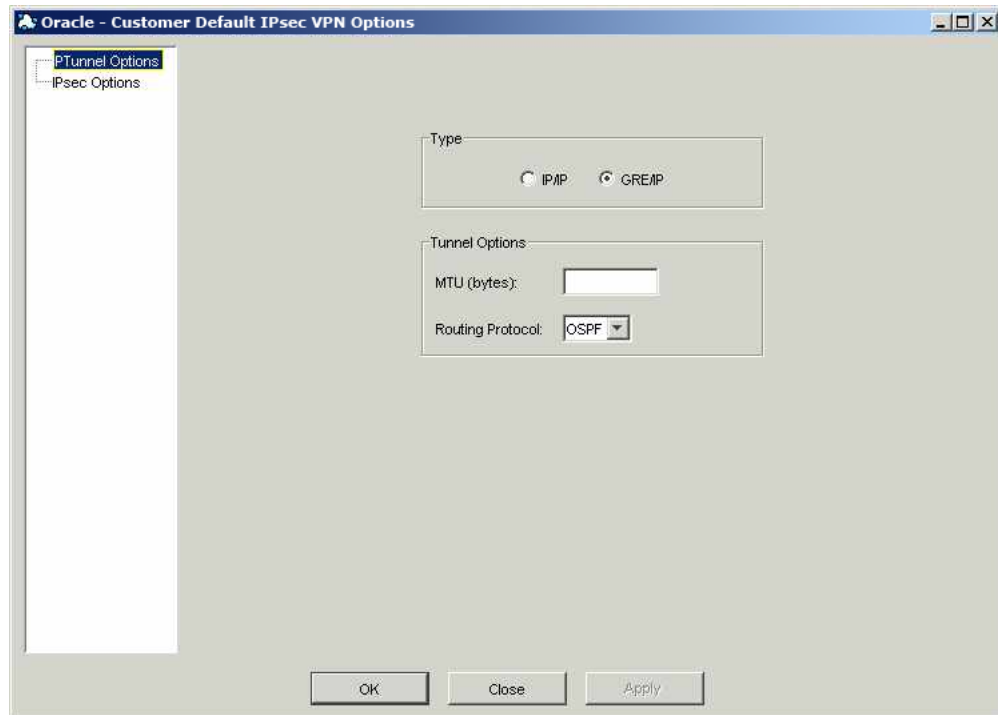
Default options are used for defaulting GUI fields when creating IPsec VPNs and IP Tunnels. If default settings are changed, it will have no effect on previously provisioned IP Tunnels and IPsec VPNs.

1. Open the Service Activator GUI and right-click on the customer for which you wish to create default options. From the drop-down, select **Modules**, then **IPsec**, then **Customer Defaults**.

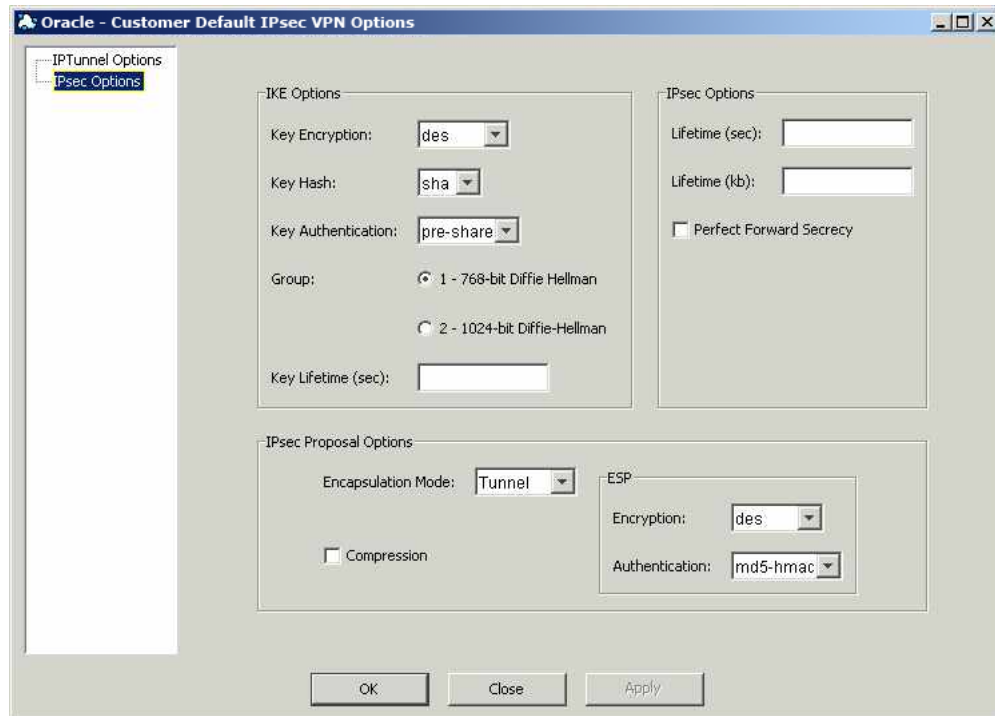


2. If this is your first time using the IPsec application, you are prompted for login information for Service Activator. See [Service Activator Login on page 116](#) for more information.

The Customer Default IPsec VPN Options dialog box appears.



3. Enter the IP Tunnel default options you wish to associate with the customer. Click on IPsec Options to display its property page.



4. Enter the IPsec default options you wish to associate with the customer. Then click 'OK'.

The button actions are as follows:

- 'Close' – exit
- 'Cancel' – exit without saving any changes
- 'Apply' – saves changes
- 'OK' – saves changes and exits

## Configuring IPsec and IP Tunnels

**Note:** To access the IPsec VPN functionality in the Service Activator GUI, ensure that the **Allow concurrent logins** option is enabled for the Service Activator user. For more information on how to do this, refer to the *Service Activator Administrator's Guide*.

### Service Activator GUI Setup

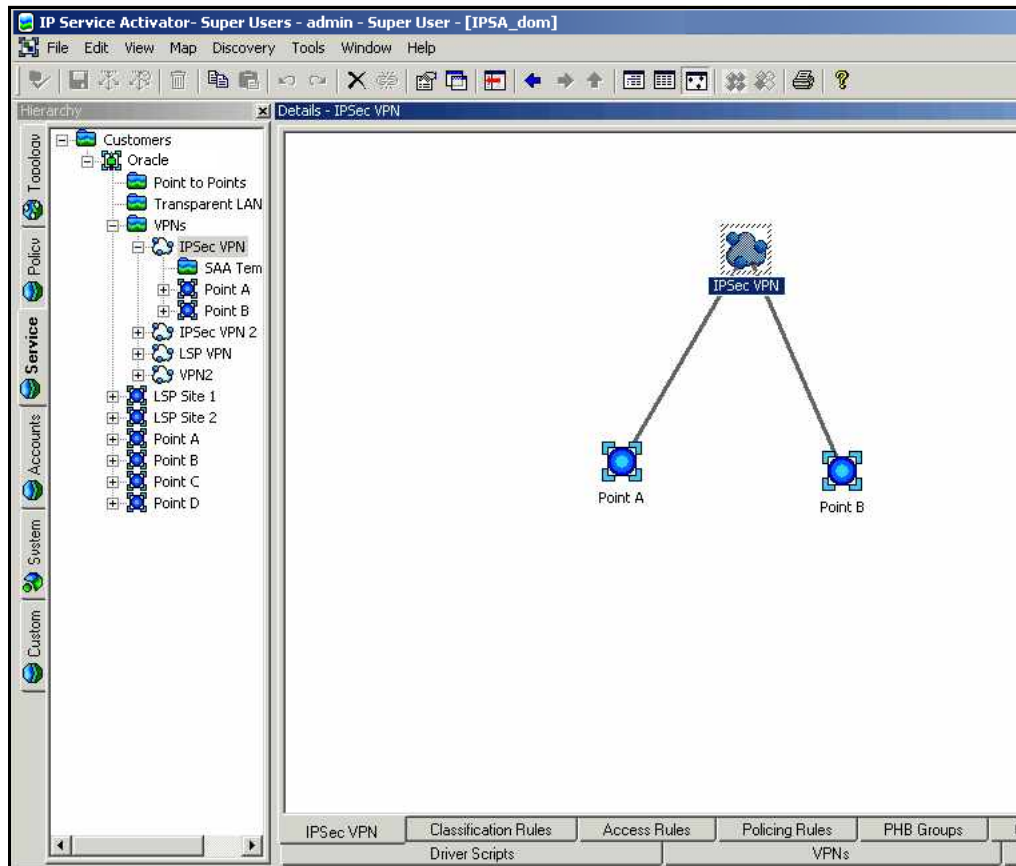
Before launching the IPsec application, the GUI needs to be pre-configured with the necessary objects to model the VPN. Perform this procedure to set up the Service Activator GUI:

1. Create a customer or use a pre-existing customer.
2. Create a VPN, fill in a name, and uncheck the MPLS checkbox (important not to miss). Leave the other parameters to the default or existing value.
3. Create two VPN sites that are the endpoints of the VPN. You only need to enter a site name; leave all other parameters to the defaults.
4. Link a discovered managed device with the role of 'Access' to each site, one device per site.



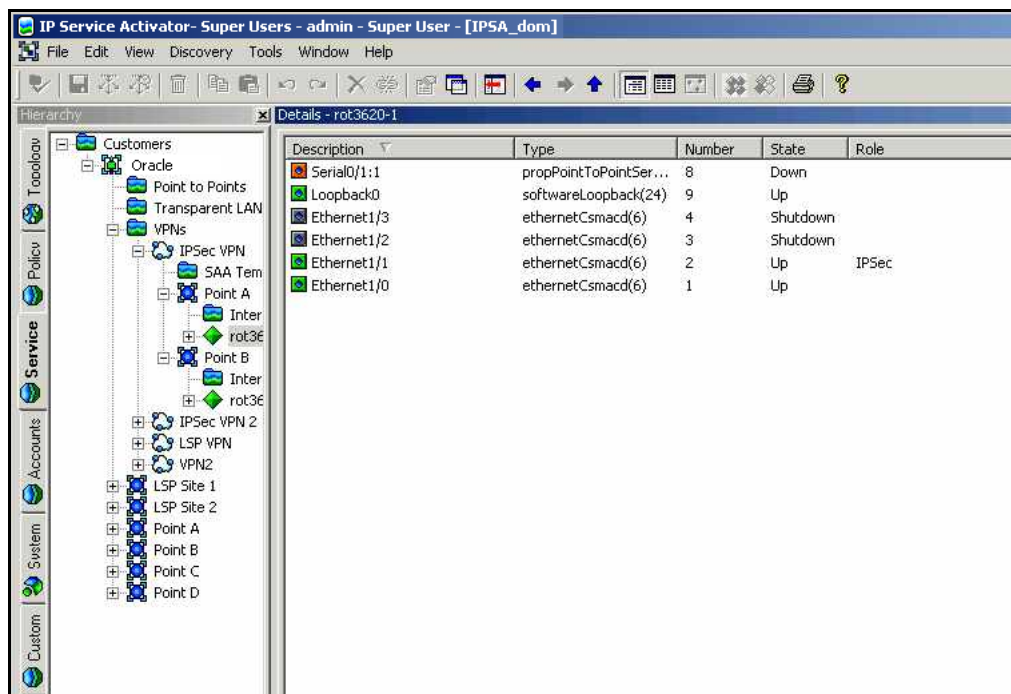
- 5. For each site, associate it with the VPN. This is done by linking the VPN with each of the sites.

See the following figure for an example setup.



- On a per device basis, assign a role of 'IPsec' to the interface(s)/sub-interface(s) to be used in the VPN. The 'IPsec' role is a custom role created during installation.

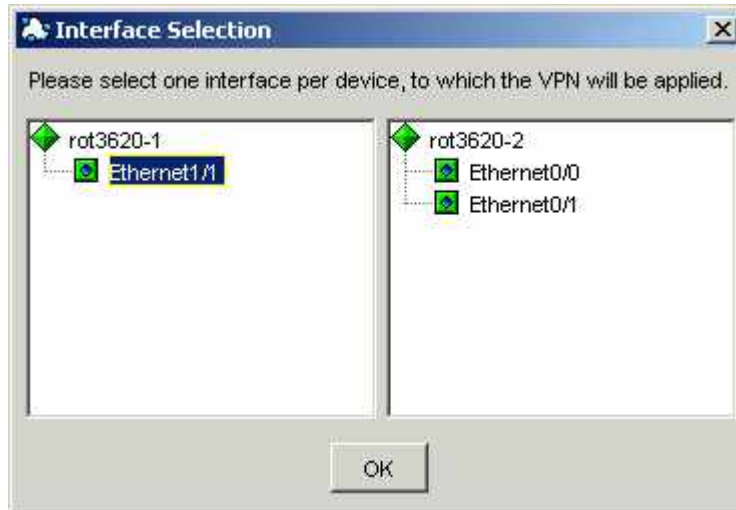
See the following figure for an example of a device's interfaces with roles of 'IPsec'.



### Multiple Interface Overview

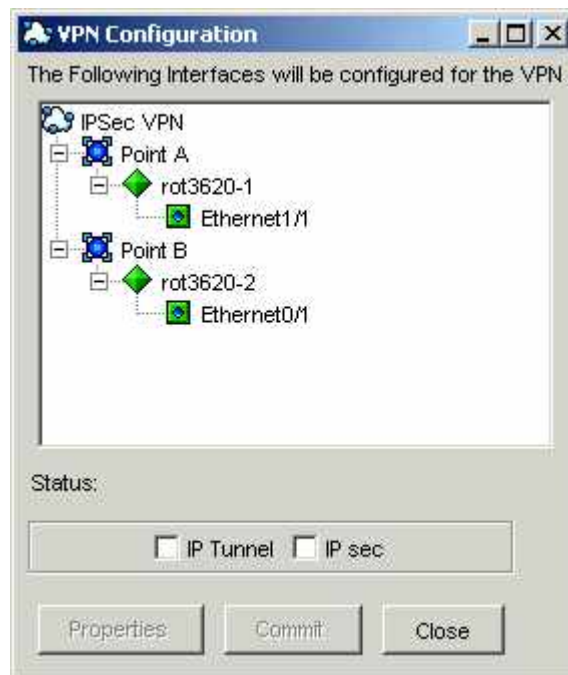
The IPsec application supports devices that have multiple interfaces or sub-interfaces with a role of 'IPsec'. To configure an IPsec VPN, the application can only use one interface per device. The Interface Selection dialog box is displayed, to pick the appropriate interface per device. If the devices in the VPN do not have multiple interfaces, this dialog box is not shown. The dialog box prevents you from picking the device or an interface that does not have a role of 'IPsec'.

The Interface Selection dialog box shows a sub-interface with a role of 'IPsec' below the interface.



## Main Summary Screen Overview

The IPsec Main Screen is the VPN Configuration dialog box, shown in the following figure. It displays a summary of relevant IPsec VPN information (device, interface, or sub-interface, and type and status of VPN being deployed). It also provides the buttons to configure, deploy and exit the application.



### GUI status field

The status field can display the following values:

- blank - no VPN is deployed.
- 'Not Installed' - you have entered IPsec or IP Tunnel information on the configuration screens but that it has not been deployed to the routers.
- 'Active' - an IPsec or IPTunnel VPN has been deployed to the routers and that there was no issues reported to our application. The information has also been stored in the database. You can view the information on the configuration screen but is not able to modify it.
- 'Rejected' - the application experienced errors creating the VPN script and/or storing the information in the database. You can modify the information and try again.

## GUI IP Tunnel and IPsec checkboxes

The IP Tunnel and IPsec checkboxes tell the application what type of VPN you wish to create and deploy or they tell you what type of VPN has been deployed. Three different VPNs can be implemented:

- IP Tunnel VPN
- IPsec VPN
- IPsec VPN over an IP Tunnel

If you have entered information for an IPsec VPN over an IP Tunnel but only want to deploy an IPsec VPN, deselect the IP Tunnel checkbox before clicking **Commit**. Note that the buttons are unchangeable when you have launched the GUI in viewing mode after the VPN has been successfully deployed.

## GUI buttons

The buttons at the bottom of the screen are used to configure a VPN (using the 'Properties' button), to deploy a VPN ('Commit' button), to re-deploy a VPN ('Re-Submit' button) and to exit the application ('Close' button). The 'Properties' and the 'Commit' buttons are enabled or disabled depending on whether the IPsec and/or IP Tunnel checkboxes are selected and whether or not the VPN has been configured/ deployed.

The **Properties** button is disabled when the status is blank and no checkbox is selected. It is enabled if you select a checkbox, or if you enter the application when you have already deployed a VPN.

The **Commit** button is disabled by default and will only be enabled if you have entered parameters in the VPN configuration screens and one of the IPsec or IP Tunnel checkboxes is selected.

The **Re-Submit** button appears in place of the **Commit** button when a VPN has been deployed and you are trying to view an existing IPsec/ IP Tunnel VPN. The purpose of this button is to resend the configuration down to the routers.

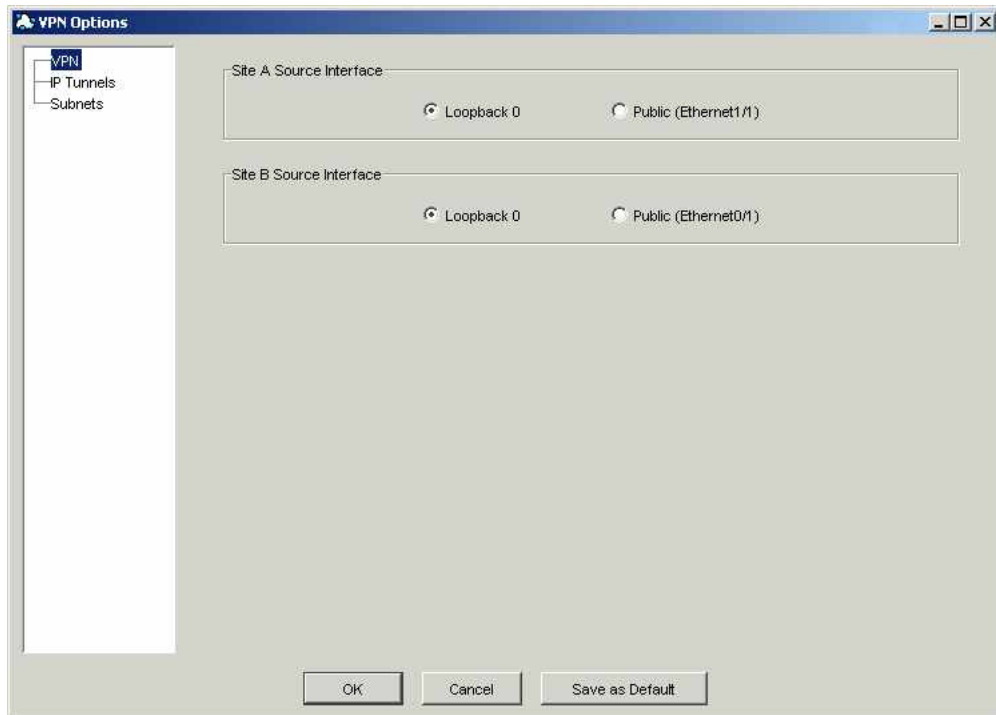
## Creating an IP Tunnel VPN

1. Set up and configure the Customer, Site, and VPN objects in Service Activator in preparation for launching the IPsec application. Perform the procedure [Service Activator GUI Setup on page 124](#) to accomplish these tasks.

The Interface Picker Screen appears if any of the devices have more than one sub-interface/interface with a role of 'IPsec'. For more information see [Multiple Interface Overview on page 126](#).

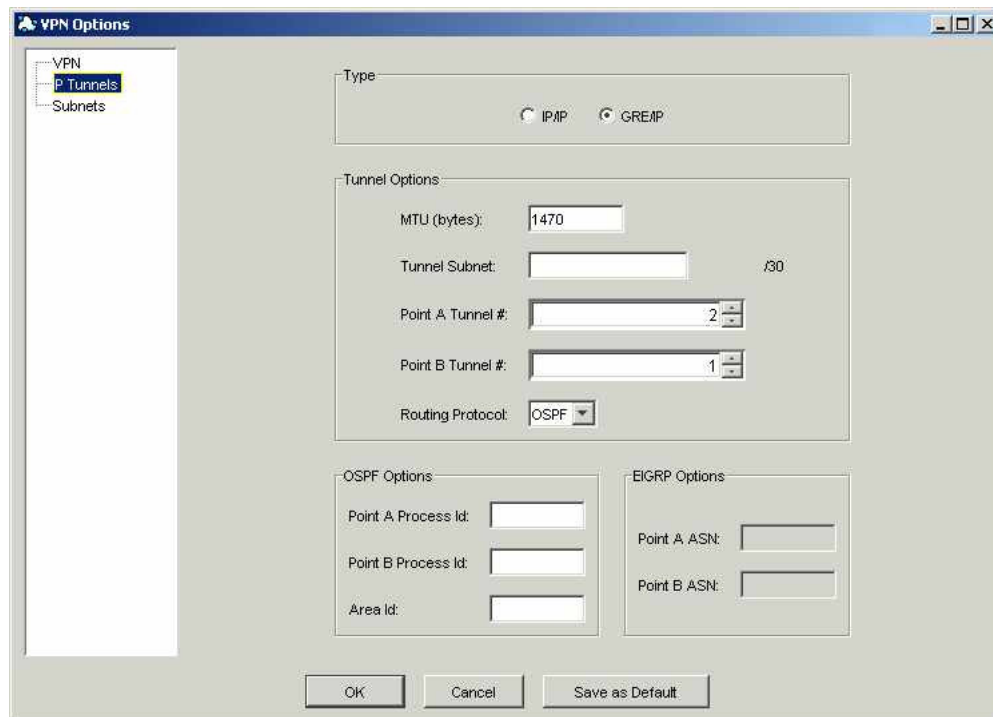
2. Open the IPsec application and display the main screen.
3. Click the IP Tunnel checkbox and click the 'Properties' button.

The VPN Options dialog box is displayed; appropriate fields are shown following.



4. Select the site's source interfaces. The source interface you pick for Site A is the peer interface for Site B and vice versa. If one of the sites is already in use by another IPsec VPN, the source interface radio button selection is grayed out.

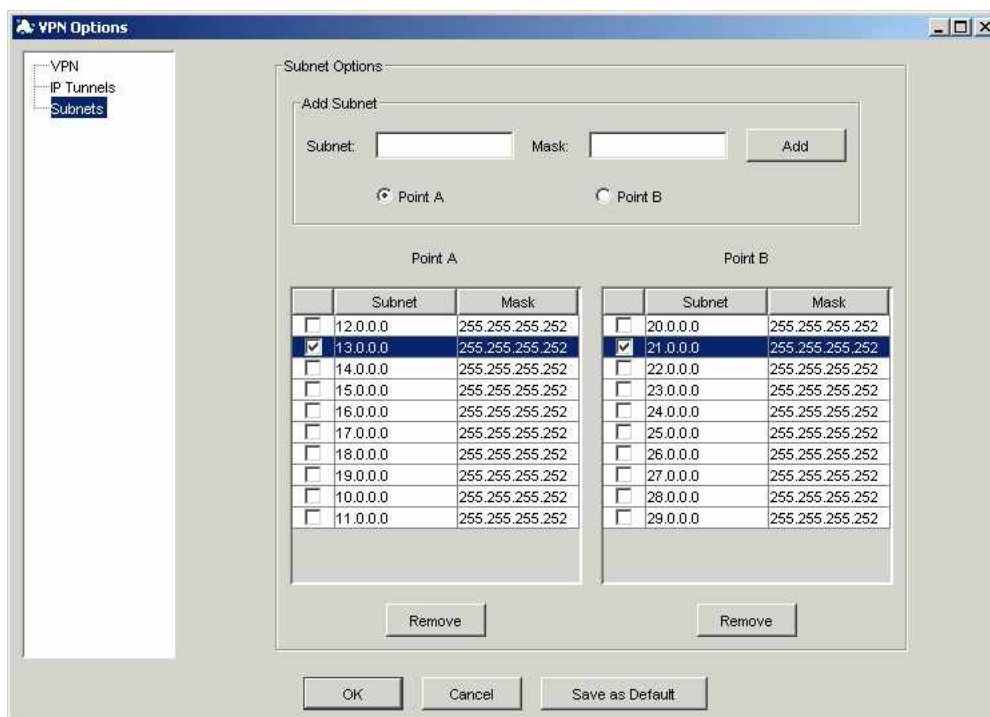
5. Click the IP Tunnels property page; appropriate fields are shown in the following figure.



The screenshot shows the 'VPN Options' dialog box with the 'IP Tunnels' property page selected. The 'Type' section has radio buttons for 'IPsec' and 'GRE/IP'. The 'Tunnel Options' section includes fields for 'MTU (bytes)' (1470), 'Tunnel Subnet' (with a '/30' suffix), 'Point A Tunnel #' (2), and 'Point B Tunnel #' (1). The 'Routing Protocol' is set to 'OSPF'. Below are sections for 'OSPF Options' (Point A Process Id, Point B Process Id, Area Id) and 'EIGRP Options' (Point A ASN, Point B ASN). At the bottom are 'OK', 'Cancel', and 'Save as Default' buttons.

6. Enter the IP Tunnel fields. Key Features of the screen are:
  - a. The mandatory fields on this screen are the tunnel type, tunnel subnet, tunnel number, the routing protocol and routing protocol options.
  - b. The tunnel number needs to be unique in the device. The EEU Database and the Service Activator object model are verified before populating the tunnel number pick box. If a tunnel number is being used, the next number in sequence is shown, that is if tunnel numbers 1,3,5,6 were being used, 2 would be displayed. The tunnel number can have different values for each device.
  - c. The tunnel subnet must be a valid /30 IP address.
  - d. The MTU can be a value between 68 and 1,000,000 bytes.
  - e. The OSPF process id can be a value between 1 and 65,535. The process id can have different values for each device.
  - f. The OSPF area id can be a value between 0 and 4,294,967,295.
  - g. The EIGRP ASN can have different values for each device.

7. Depending on which routing protocol is selected, perform a, b, or c sub-step:
  - a. If OSPF is selected, the OSPF options process id and area id become editable and need to be filled in.
  - b. If static routing is selected, the subnets property page becomes enabled and subnets need to be entered.
    - i. Click the Subnets property page; appropriate fields are shown in the following figure.



- ii. Enter a subnet and mask, a minimum one per site.
- iii. Click the 'Add' button to put the subnet and subnet mask in the list. Enter as many subnets as necessary. A selected checkbox next to the subnet indicates that this subnet will be used in the VPN.
- iv. Click the radio button next to your other site name, and perform step iii) as needed.
- v. To remove a subnet, select/highlight the subnet and mask from the list and click the 'Remove' button.
- c. If EIGRP routing protocol is selected, enter a value for both EIGRP Options ASN field. These values can be different since they are applied per device.



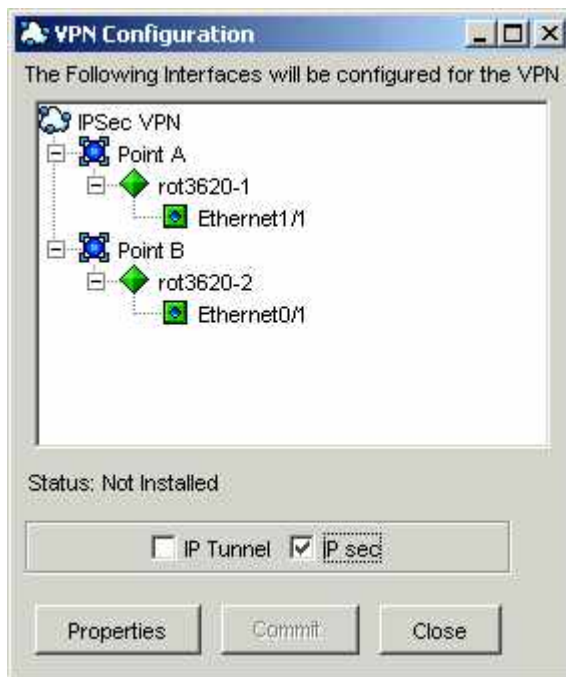
8. If you want to overwrite the customer defaults, click the 'Save as Default' button.
9. Click 'OK' to save entered values and return to the main screen (or click 'Cancel' if you want to return to the main screen without saving entered values).
10. To deploy the IP Tunnel VPN on the router, click the 'Commit' button.
11. It is recommended to clean up driver scripts after the provisioning process is complete. See [Cleaning up Driver Scripts on page 141](#) for details.

## Creating an IPsec Tunnel VPN

1. Set up and configure the Customer, Site, and VPN objects in Service Activator in preparation for launching the IPsec application. Perform the procedure [Service Activator GUI Setup on page 124](#) to accomplish these tasks.

The Interface Picker Screen appears if any of the devices have more than one sub-interface/interface with a role of 'IPsec'. For more information see [Multiple Interface Overview on page 126](#).

2. Open the IPsec application and display the main screen.

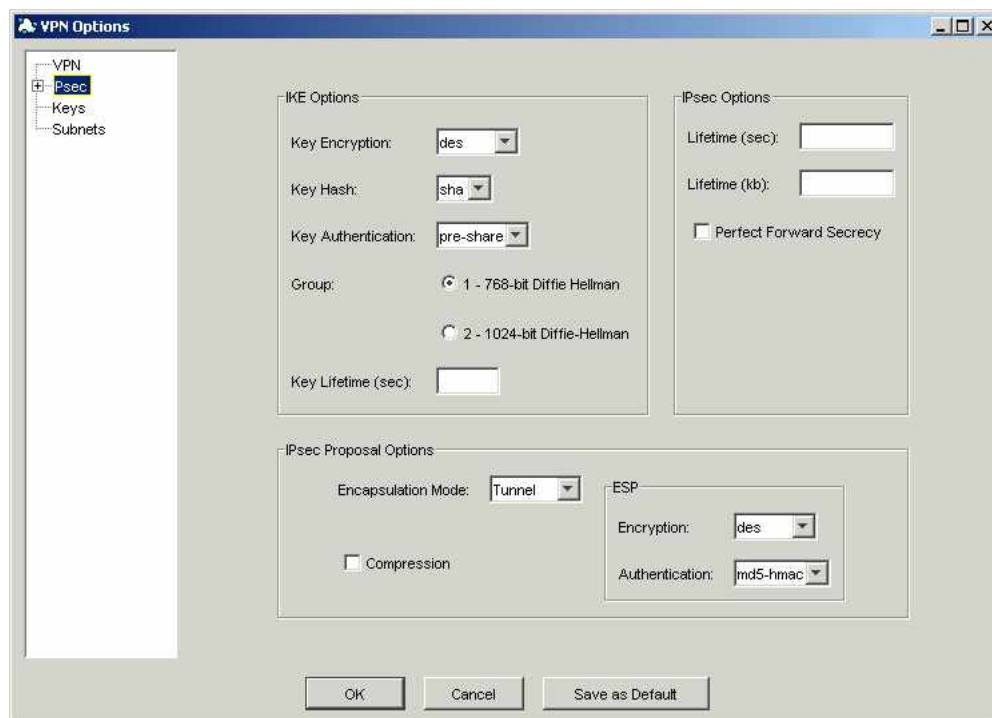


3. Click the IPsec checkbox and then click the 'Properties' button.

The VPN Options window opens.

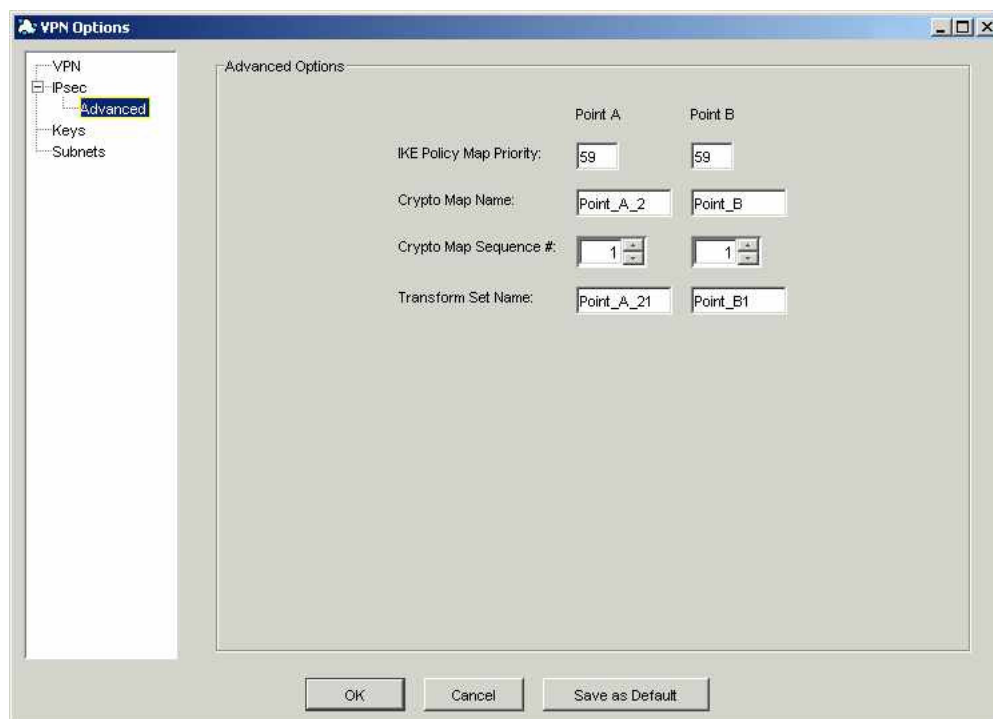


4. Select the site's source interfaces. The source interface you pick for Site A is the peer interface for Site B and vice versa. If one of the sites is already in use by another IPsec VPN the source interface radio button selection will be grayed out

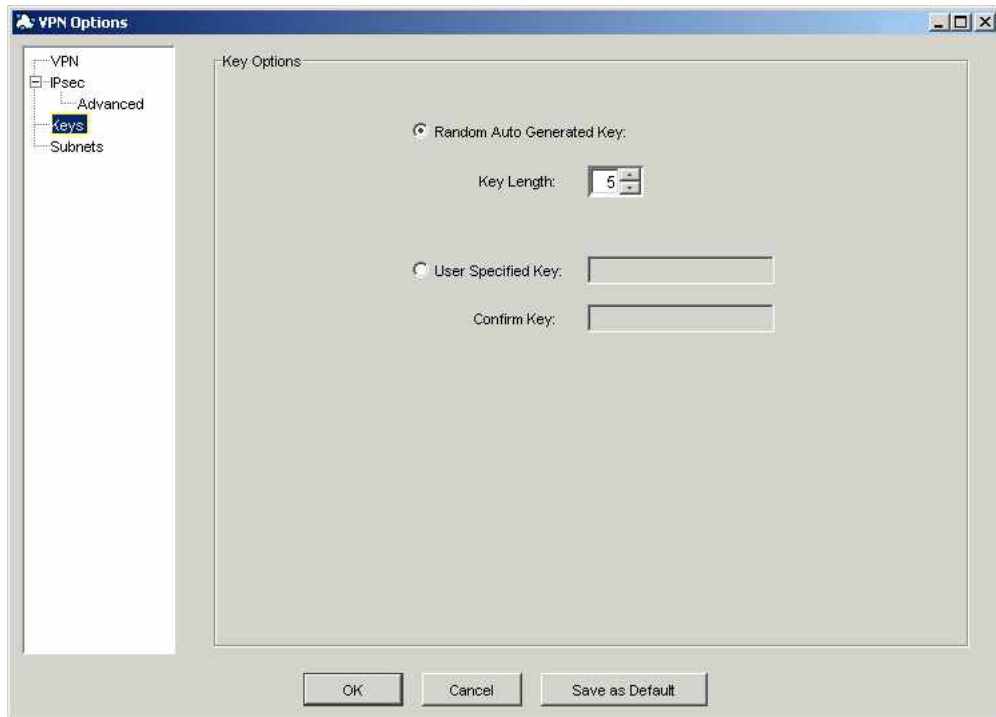


5. Select the IPsec property page. Enter the IPsec Options. Key features of this screen are:
  - Fields will be pre-populated with Customer or Global default options, if they exist.
  - Key Lifetime (sec), Lifetime (sec), and Lifetime (kb) are optional fields.
  - Lifetime (sec) must be a value between 120 and 86,400.
  - Lifetime (kb) must be a value between 2,560 and 536,870,912.
  - Key Lifetime (sec) must be a value between 43,200 and 172,800.
6. To save these IPsec options as Customer Default options, click the 'Save as Default' button. Future IPsec VPNs for this customer will default to the same values as the current one.
7. Select the '+' next to the IPsec property page to override the defaults for Advanced property page. Key features of this screen are:

- Point A Crypto Map Name and Point B Crypto Map Name are defaulted to the site name for device A and site name for device B respectively. If the site name is already used as a crypto map name on the device, then an underscore and a sequential number are appended (ex. \_1, \_2...etc). You can change the default crypto map name as long as the value is unique on the device.
- If the interface that is used for this IPsec VPN already has a crypto map name assigned to it, that value cannot be changed and therefore the Point A Map Name and/or Point B Map Name field are un-editable.
- The Transform Set Names (Point A and Point B) are defaulted to a combination of the Map Names plus the appropriate Point Crypto Map Sequence Number.
- IKE Policy Map Priority is defaulted with a system-generated sequence number. It is recommended that the value not be changed (but can be if required). A different value can be entered for each of Point A and Point B.
- Crypto Map Sequence Number can have both the same values or a different value can be entered for each of Point A and Point B. The values are verified with what sequence numbers are already in use.

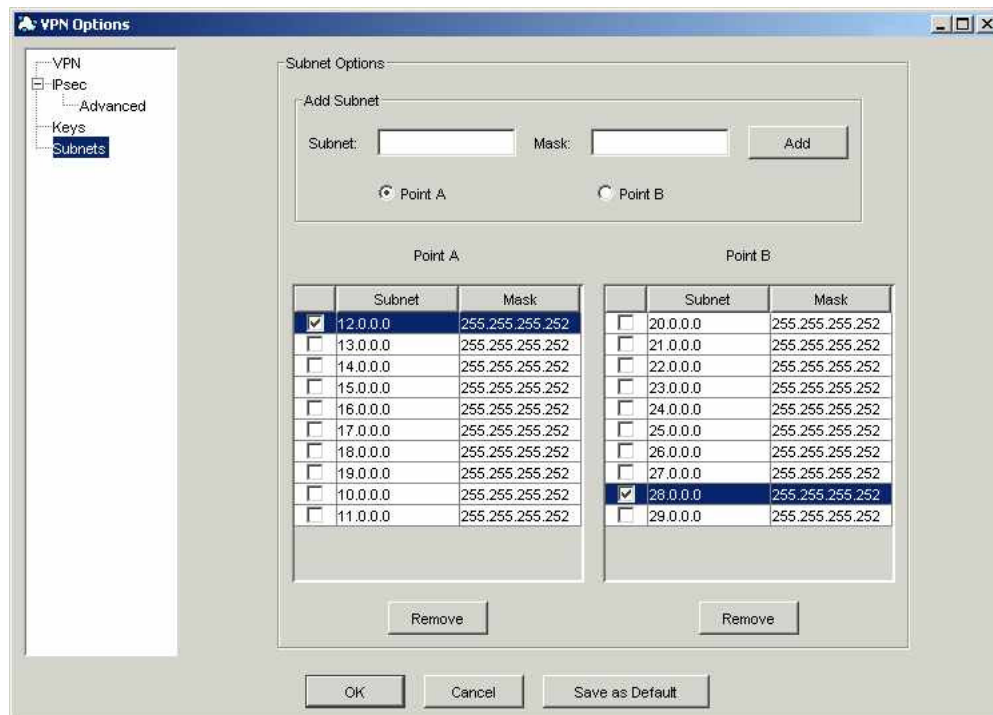


8. Select the Keys property page; appropriate fields are shown in the following figure.



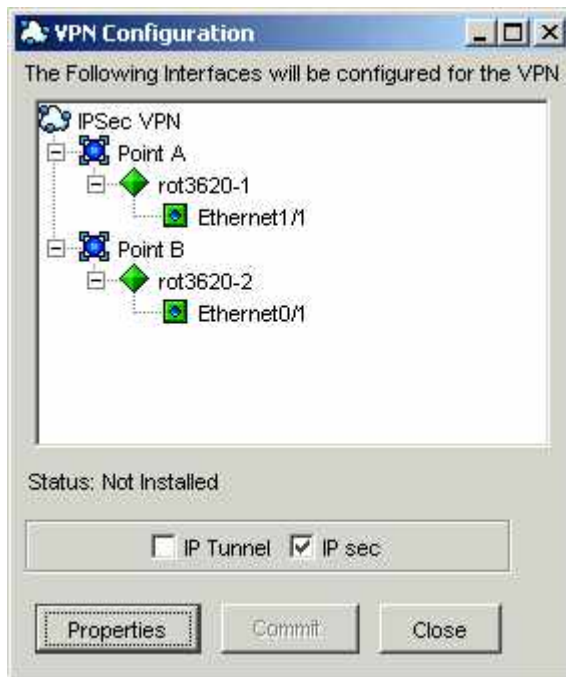
9. By default, the pre-shared key is a randomly generated value with a length of 5 characters. The length can be changed to a maximum of 50 characters. The User Specified Key can also be selected, in which case the key must be entered twice to confirm the value.

10. Select the Subnets property page; appropriate fields are shown in the following figure.



11. Enter the subnets for this IPsec VPN. The names of the two sites involved in the VPN are displayed above the tables.
- a. Enter a subnet and mask, a minimum of one per site
  - b. Click the 'Add' button to put the subnet and subnet mask in the list below. Enter as many subnets as necessary. A selected checkbox next to the subnet indicates that this subnet will be used in the VPN.
  - c. Click the radio button next to your other site name, and perform step ii) as needed.
  - d. To remove a subnet, select/highlight the subnet and mask from the list and click the 'Remove' button.
12. When you are finished entering all options for the IPsec VPN, click the 'OK' button.

The IPsec Main Screen is displayed.



13. To apply the configuration to the routers, select the 'Commit' button. The application will close the window.

## Creating an IPsec over IP Tunnel VPN

To set up and configure an IPsec over IP Tunnel VPN select both IP Tunnel and IPsec checkboxes. The screens are identical as if you were configuring IPsec or IP Tunnel individually. Refer to [Creating an IP Tunnel VPN on page 129](#) and [Creating an IPsec Tunnel VPN on page 133](#) for specific configuration details.

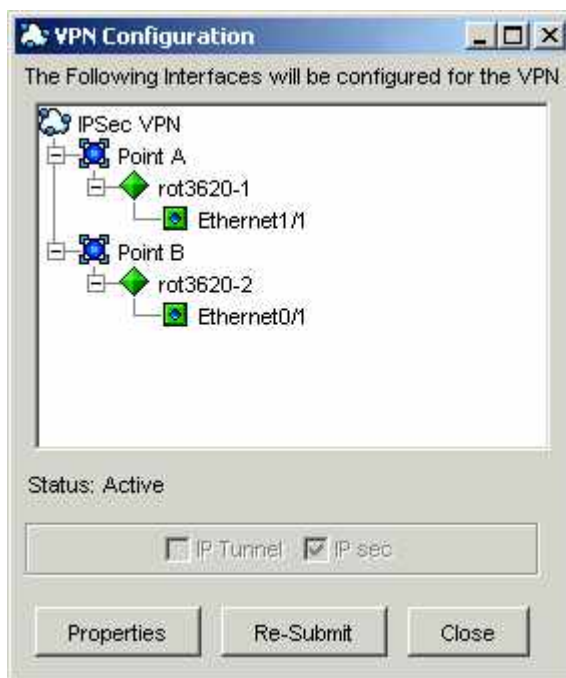
**Note:** When configuring an IPsec over IP Tunnel VPN, the 'Subnets' property page is associated with the routing of the IP tunnel. Therefore, the 'Subnets' property page will only be available when 'Static' is selected as the IP tunnel routing type.

## Administration Tasks

### Viewing an Existing IPsec/ IP Tunnel VPN

Use this procedure after you have created a VPN using the [Creating an IPsec Tunnel VPN on page 133](#) procedure.

1. Launch the IPsec Application by right-clicking the mouse on the VPN icon and selecting Modules -> IPsec -> Create.



The appropriate checkbox, depending on the type of VPN, is checked but disabled.

2. Click the 'Properties' button to viewing the pre-configured VPN Options.

The VPN Options screen is shown with all fields disabled.

If an error was raised in the fault pane of Service Activator and you have addressed the reason for the fault you can re-submit the VPN configurations down to the routers.




## Deleting an IPsec/IP Tunnel VPN

**Important Note:** Do not unlink or delete VPN, Site, or Device objects that are used in IPsec/IP Tunnel VPNs, out of Service Activator until steps 1 and 2 below have been completed first.

1. Launch the IPsec Application by right-clicking the mouse on the VPN icon and selecting Modules -> IPsec -> Delete.

The VPN\_Delete Screen appears.



2. Click the 'OK' button to confirm that you wish to delete the VPN, or click the 'Close' button if you do not want to delete the VPN.
3. It is recommended to clean up driver scripts after the delete process is complete. See [Cleaning up Driver Scripts](#) for details.

Upon deletion of the VPN the router configuration is removed from the router and the IPsec Options are removed from the database.

## Cleaning up Driver Scripts

When an IP/GRE Tunnel or IPsec VPN Tunnel is provisioned, driver scripts are attached to both devices involved in the point-to-point connection. These driver scripts, which hold data relevant to the tunnel creation, are used only once. They can be removed after the driver script state changes to 'Installed'.

The name of the script is a combination of the application name (IPsec), the operation (Create or Delete), the device name and the time value in milliseconds. For example, if a VPN was created on a device called "rot3640-1", the script name would be "IPsec\_Create\_rot3640-1\_235665544.py".

To clean up the driver scripts, refer to the Datascript deletion procedure in the *Administrator's Guide*.

## VRF-Aware IPsec

The VRF-Aware IPsec feature allows you to map IPsec tunnels that terminate on a shared public interface to specific Virtual Routing and Forwarding (VRF) instances, therefore allowing you to map IPsec tunnels to MPLS VPNs. This allows you to extend customer VPN access to users that are not directly reachable via dedicated WAN links.

**Note:** For details on implementing VRF-Aware IPsec using the Service Activator including step-by-step procedures, refer to the Service Activator Online Help.

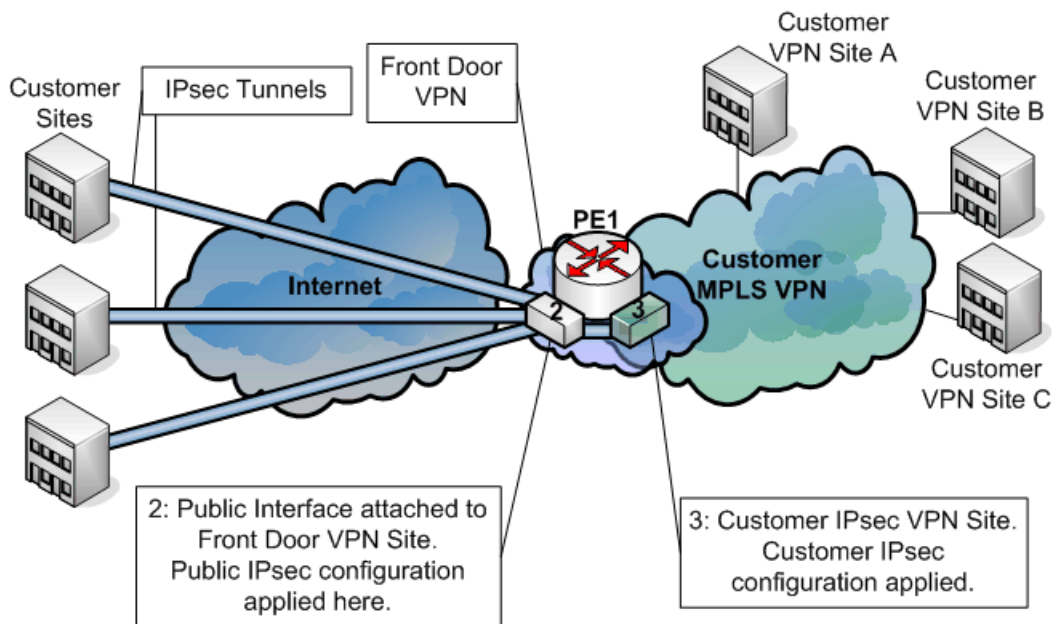
### Example VRF-Aware IPsec implementation

To support the VRF-Aware IPsec feature, two VRF instances are used: the Front Door VRF (F-VRF) at the Front Door Site and Customer VRF (C-VRF) at the Customer IPsec Site. The Front Door VRF is used to isolate the public interface from the global routing tables on the PE. The C-VRF provides connectivity to the Customer MPLS VPNs.

One or more IPsec tunnels can be attached to the PE using a single public interface. The tunnels then terminate at the appropriate C-VRF. The C-VRF for each of the IPsec tunnels may be different depending on the VRF referenced from the customer specific IPsec configuration.

#### **Service Activator's VRF-Aware IPsec provisioning support includes:**

- for the MPLS VPN portion, provisioning and activation of customer VPN Sites and VPNs.
- for the IPsec portion, provisioning and activation of IPsec configuration policies and their association to the Front Door Site and specific Customer MPLS VPN sites



The diagram above illustrates a scenario where there is an existing customer VPN containing various sites (Customer MPLS VPN Sites A, B and C). Connectivity between the Remote sites and the Customer MPLS VPN sites is desired. However the Remote sites are not directly attached to the WAN supporting the Customer MPLS VPN.

IPsec tunnels are configured between the Remote Sites and the PE device. They all connect to the PE through the same public interface (labelled 2 on the diagram). This interface is attached to the Front Door Site (labelled 2) in the Front Door VPN. Public IPsec configuration is applied to the Front Door Site.

The Customer IPsec Site (labelled 3) is a member of both the Front Door VPN and the Customer MPLS VPN. The PE device containing the public interface is associated with Customer IPsec Site. The customer IPsec configuration policy is applied to the Customer Site.

Traffic moving from IPsec Sites to Customer VPN Sites is encrypted at the start of the IPsec tunnels. Packets travel from the IPsec Sites over a publicly accessible network such as the Internet, through the IPsec tunnels, through the public interface (2) on the PE. They are passed to the Customer IPsec site where they are decrypted.

The appropriate MPLS labelling is applied for transport over the Customer MPLS VPN. Going in the other direction, MPLS-labelled packets arrive at the Customer

IPsec site. The MPLS labels are stripped and the packets are encrypted for transport over the IPsec tunnels. At the IPsec Sites, they are decrypted.

## **Appendix A**

# **Setting Up Management and Customer VPNs**

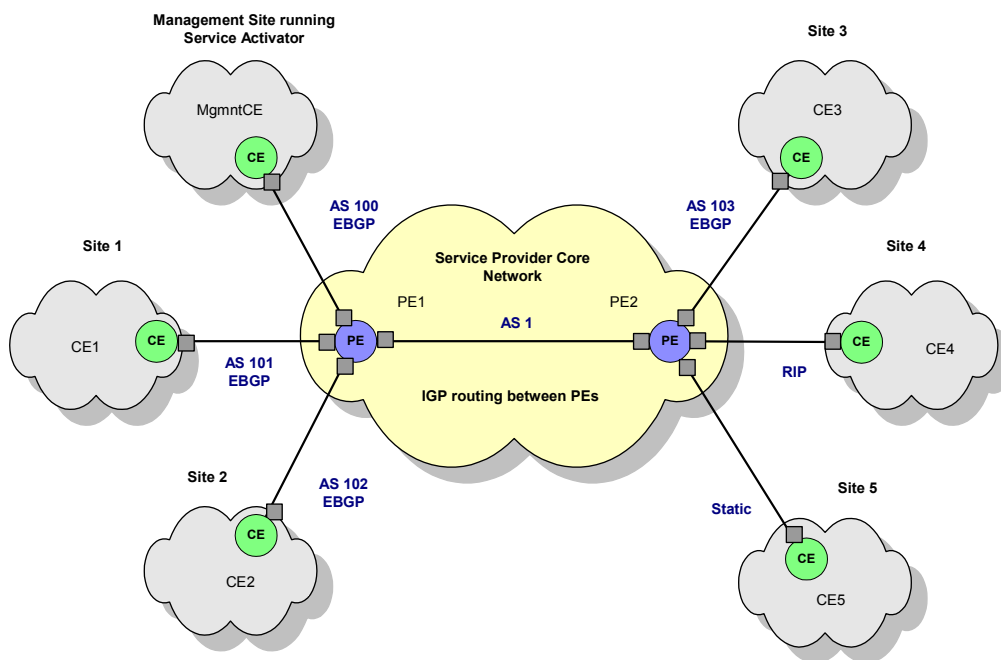
This appendix outlines the high-level steps you need to follow in order to set up management and customer VPNs correctly.

## Introduction

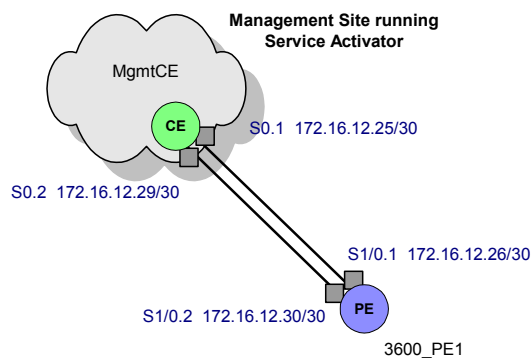
In the following example, three VPNs are created:

- A management VPN is set up comprising the Management site and all other sites in all VPNs managed by Service Activator. With Service Activator running at the management site, all CE routers can be managed. The management site is a hub site; all other sites are spoke sites.
- Customer VPN 1 comprises sites 1, 2 and 4.
- Customer VPN 2 comprises sites 1, 3 and 5.

VPN 1 and 2 are both fully-meshed. The network layout and routing protocols are shown in the following diagram:



Note that between the management site and the core network, two links are required, one to provide VPN connectivity and one to provide routes to the Service Provider backbone IGP. For example:



## Steps to configure the VPNs

The following steps explain the sequence of operations required to set up a management VPN, enabling access to the CE routers, followed by customer VPNs. For full details, see the detailed explanation in [Setting Up RFC2547 MPLS VPNs on page 1](#).

### 1. Set up domain

Firstly, set the domain properties. On the **Domain** page of the Domain dialog box, the **Type** is set to **MPLS VPN**, and the **Public PE to CE Addresses** checkbox is selected.

Set the required parameters on the **VPN BGP** page and the **VPN MPLS** page. See [Setting up domain parameters on page 15](#).

### 2. Set up core network ASN

On the **ASN** page of the **Domain** dialog box, set the internal BGP ASN to 1. See [Setting up the provider core ASN on page 18](#).

### 3. Discover the network and assign roles to devices and interfaces

Run a device discovery to discover the PE devices (PE1 and PE2) and the CE device at the management site (MgmtCE). They can be discovered using their IP addresses or DNS names.

Assign the correct system-defined roles to devices, i.e. PE routers must be assigned the Gateway role and the CE router must be assigned the Access role. Appropriate

interfaces on the devices also need to be assigned Local or Access roles. Role assignment is normally performed automatically by means of user-defined role assignment rules, otherwise you need to manually assign a role to each device and interface.

#### 4. Disable interfaces on CE - PE link

The interfaces at both ends of the Serial 0.1 link between PE1 and MgmtCE, i.e. the link that is not to be used for the VPN connection, must be assigned a role of **Disabled**, to prevent them being configured into a VPN.

Note also that the interface or sub-interface that is in the Management VPN and thus provides routes to the CE must be passive. This is so that routes from the customer networks are not leaked into the Service Provider backbone, and vice versa.

#### 5. Set devices to Managed

Set all discovered devices to **Managed** so that Service Activator can configure them.

#### 6. Set up customers and sites

Set up appropriate customers on the **Service** tab. Create a dummy customer, such as "Management" for the management VPN, and "Customer 1" and "Customer 2" for the customer VPNs. Create site objects for the management site and each of the customer sites and give them identifying names.

#### 7. Link physical network components with sites

For each site, link the appropriate access interface on the PE router to the site by dragging and dropping.

#### 8. Set the VPN routing parameters

For each site, set the appropriate routing type and relevant parameters (EBGP, RIP, OSPF and/or static routing) on the **Site** property pages.

For each site, you need to set up private addresses (used within the VPN) and public addresses (used outside the VPN). Set these on the **Addressing** page on the **Site** dialog box.

In this example, it is assumed that the settings on the **Advanced VPN** page are left as defaults.

#### 9. Create the management VPN

Create a VPN object under the "Management" customer object to represent the management VPN.



Set up the management VPN by linking all the sites (Management site, Site 1, Site 2, Site 3, Site 4 and Site 5) to the VPN object by dragging and dropping.

On the **Connectivity** page of the VPN properties, set the Connectivity to **Management** and ensure the management site is selected as the hub site.

### **10. Implement the management VPN**

Save the changes to the database and commit the transaction to set up the management VPN and configure devices.

### **11. Discover the CE devices**

The CE devices can now be discovered, using their defined loopback addresses. Link the CE devices to the appropriate sites by dragging and dropping (note that an error is flagged if this is not done).

### **12. Create the customer VPNs**

Create two VPN objects on the **Service** tab to represent the two customer VPNs.

Set up the customer VPNs by linking the customer sites to the appropriate VPN objects by dragging and dropping:

- Customer VPN 1: Site 1, Site 2 and Site 4.
- Customer VPN 2: Site 1, Site 3, Site 5.

### **13. Implement the customer VPNs**

Save the changes to the database and commit the transaction to set up the VPNs and configure devices.

For example VPN configurations, see the appropriate Device Driver Guides for Cisco, Juniper and Unisphere.

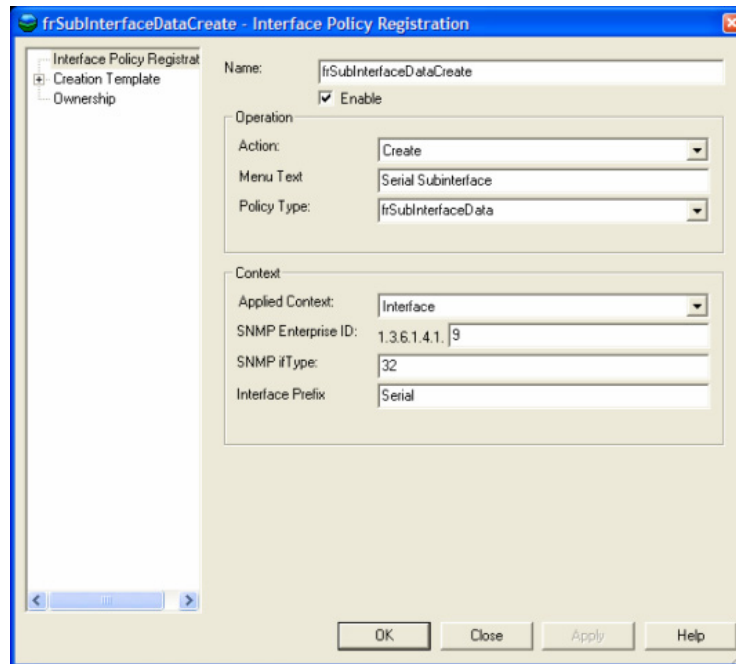


## **Appendix B**

# **Example Usage: Interface Configuration Management Module**

This appendix outlines some practical examples of using interface management configuration policies in the Interface Configuration Management Module.





3. Go to **Creation Template** tab and set the additional attributes:
  - **SNMP ifType**: 32 (frameRelay)
  - **Speed**: 0
  - **Interface Name Prefix**: Select **Get prefix from parent**
4. Click **OK** to commit the transaction.

**Note:** The Speed attribute is used to set the initial Interface speed value in the Object Model, but is not used for configuration of the device/interface—The object model value is automatically updated when the interface is discovered.

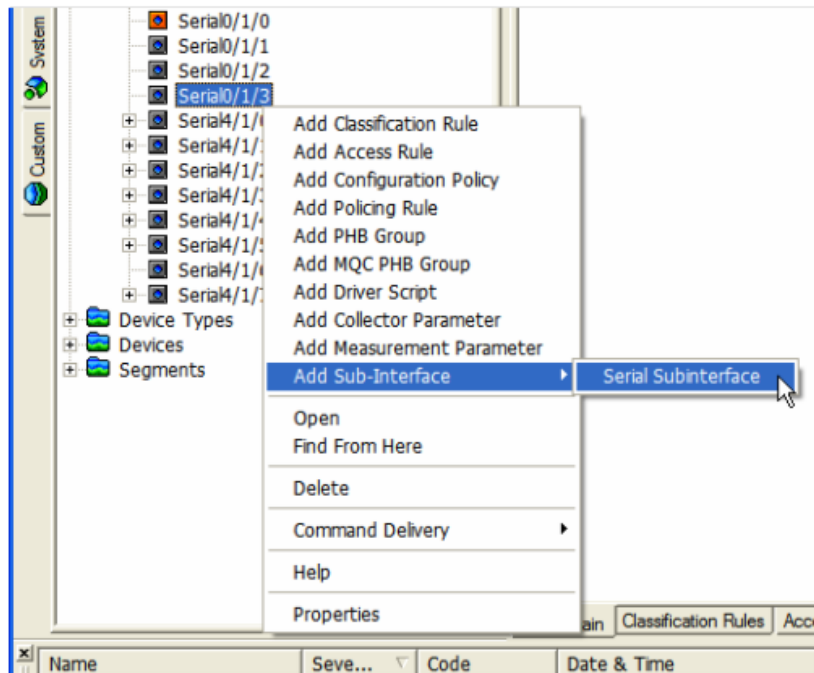
## Creating a Sub-interface

Follow these steps to create a sub-interface:

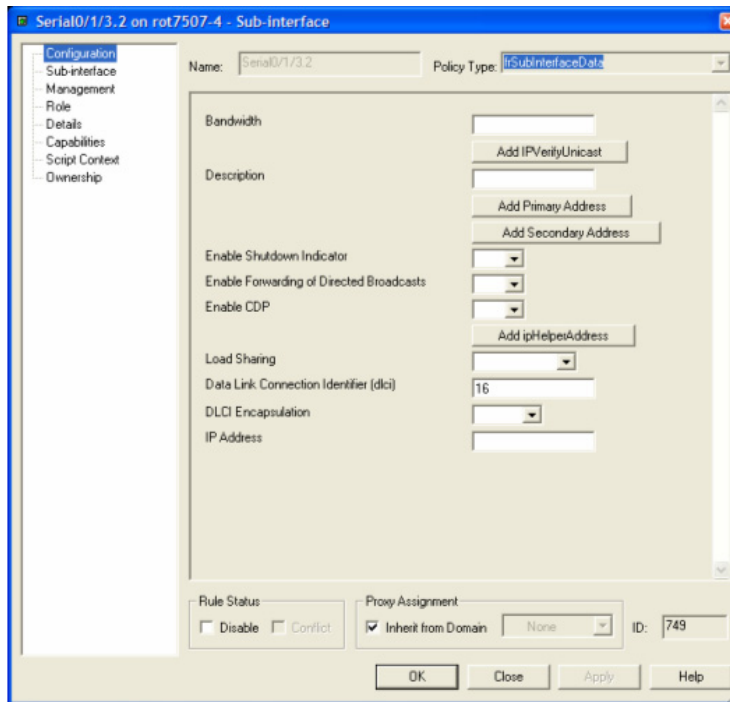
1. For a managed Device, set the appropriate Device and Interface roles. Right click a Serial interface and select **Add Sub-interface > Serial Subinterface**. This option is created by newly registered Interface Creation Policy.

**Note:** The encapsulation of the parent interface must be set to Frame Relay—`ifType frameRelay(32)`. The Add Sub-interface menu will not appear if the Serial interface type is `prepPointToPointSerial(22)` or `ppp(23)`. To change the

interface encapsulation configure the parent interface with encapsulation frame-relay and rediscover the device.



2. The interface properties dialog appears with the contents of the frSubinterfaceData configuration policy. Set the sub-interface data as appropriate.
  - Name: *Serial0/1/3.2* (mandatory field)
  - Data Link Connection Identifier (dci): *16* (mandatory field)



3. Click the **Role** tab and set the **Interface role** as **Appropriate**. If the role is not set, the interface will not be created.
4. Click **OK** to Commit the transaction.

The following Serial sub-interface is created on the device.

```
2007-09-06 20:02:14|10.156.68.126|interface Serial0/1/3.2 point-to-point
2007-09-06 20:02:14|10.156.68.126|frame-relay interface-dlci 16
2007-09-06 20:02:14|10.156.68.126|interface Serial0/1/3.2 point-to-point
2007-09-06 20:02:14|10.156.68.126|exit
```

## Removing a Sub-interface

If the sub-interface object is deleted from the IPSA object mode, the sub-interface on the device will also be removed.

```
2007-09-06 20:05:19|10.156.68.126|interface Serial0/1/3.2 point-to-point
2007-09-06 20:05:19|10.156.68.126|no frame-relay interface-dlci 16
```

```
2007-09-06 20:05:19|10.156.68.126|exit
2007-09-06 20:05:19|10.156.68.126|no interface Serial0/1/3.2
```

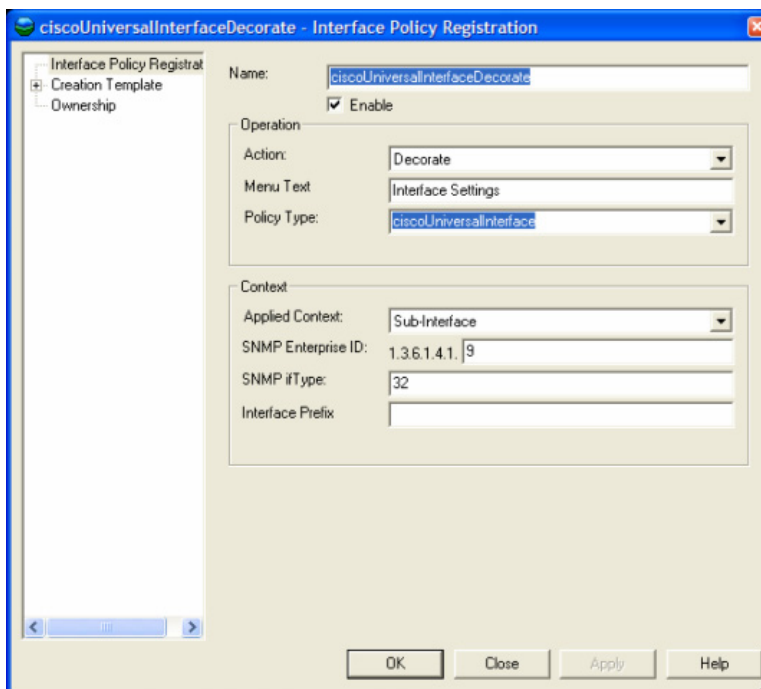
## Example: Cisco Serial Sub-interface Decoration

In this example, a Cisco Serial Sub-interface is decorated

### Interface Creation Policy Registration

To perform an interface creation policy registration, do the following:

1. On the **Policy** tab right click the **Interface Policy Registration** folder and select **Add Interface Registration**.



2. In the Interface Policy Registration dialog set the following attributes:
  - **Name:** ciscoUniversalInterfaceDecorateFrameRelay
  - **Action:** Decorate
  - **Menu Text:** Interface Settings
  - **Policy Type:** ciscoUniversalInterface

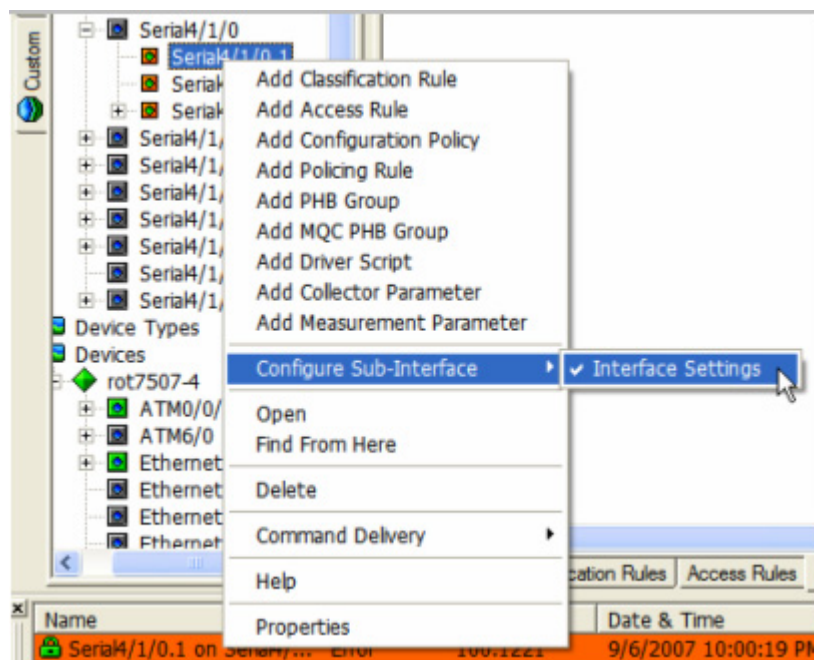


- **Context:** Sub-Interface
  - **SNMP Enterprise ID:** 9 (Cisco)
  - **SNMP ifType:** 32 (frameRelay)
3. Click **OK** to Commit the transaction.

## Decorate an existing sub-interface

Follow these steps to decorate a sub-interface:

1. For a managed device set the appropriate Device and Interface roles. Right click a Serial sub-interface and select **Configure Sub-interface > Interface Settings** option created by newly registered Interface Decoration Policy.



2. The interface properties dialog will appear with the contents of the `ciscoUniversalInterface` configuration policy. Set the sub-interface data as appropriate.

## Removing Sub-interface decoration

To remove a sub-interface decoration, right click the same Serial sub-interface and re-select the **Configure Sub-interface > Interface Settings** option created, uncheck and remove the interface decoration. The decorated settings will be removed from the interface.

**Note:** Deleting a decorated interface from the IPSA object mode will remove the decorated configuration from the device but will not remove the interface. The interface will be recreated in the object model the next time the device is re-discovered.

## Example: Cisco Channelized Serial Interface Creation

This section shows an example of Cisco Channelized Serial Interface Creation.

### Pre-requisites

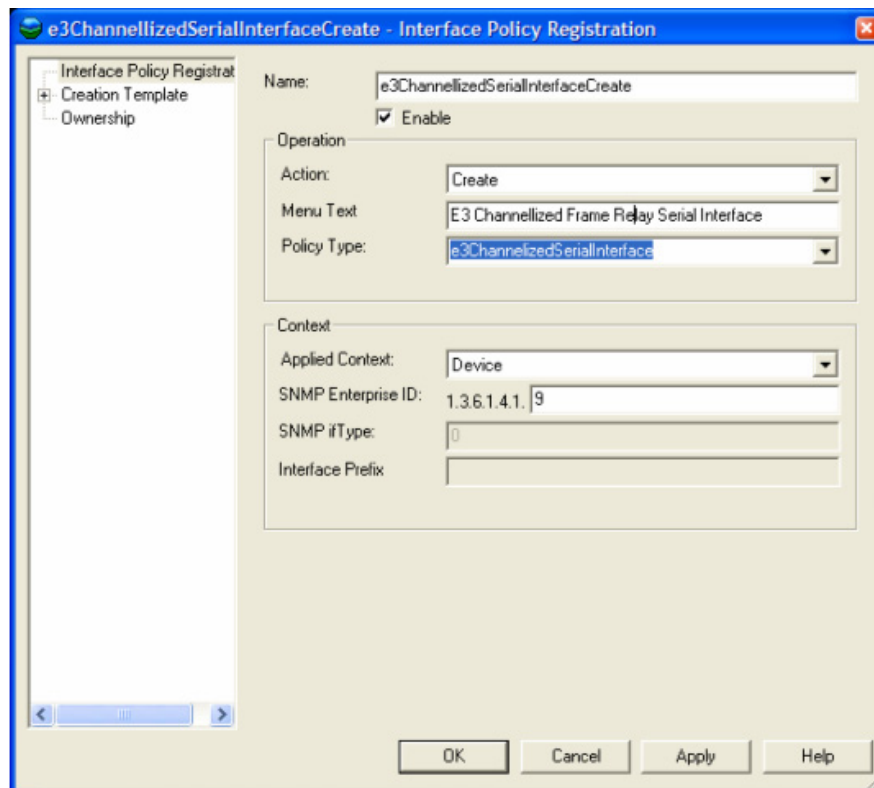
1. Modify the `Config/AutoDiscovery.cfg` file to enable Controller discovery
2. Follow the Controller Discovery comments in the `AutoDiscovery.cfg` file to comment and uncomment the appropriate lines
3. Restart the Policy Server
4. Discover the Device

**Note:** The controllers are discovered on the device as a special interface type with a different icon.

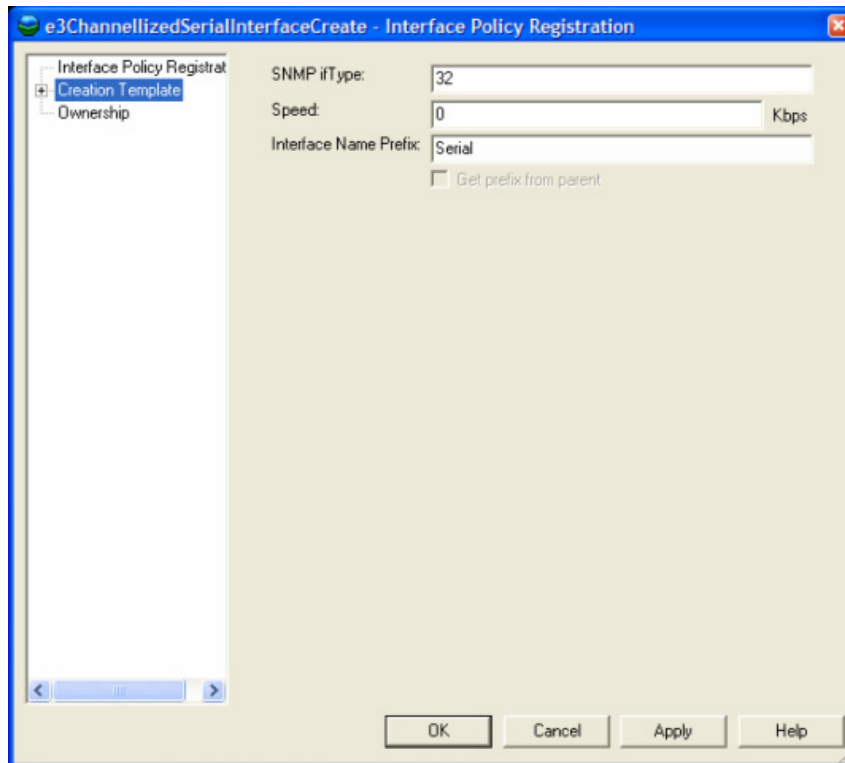
### Interface Creation Policy Registration

To perform an interface creation policy registration, do the following:

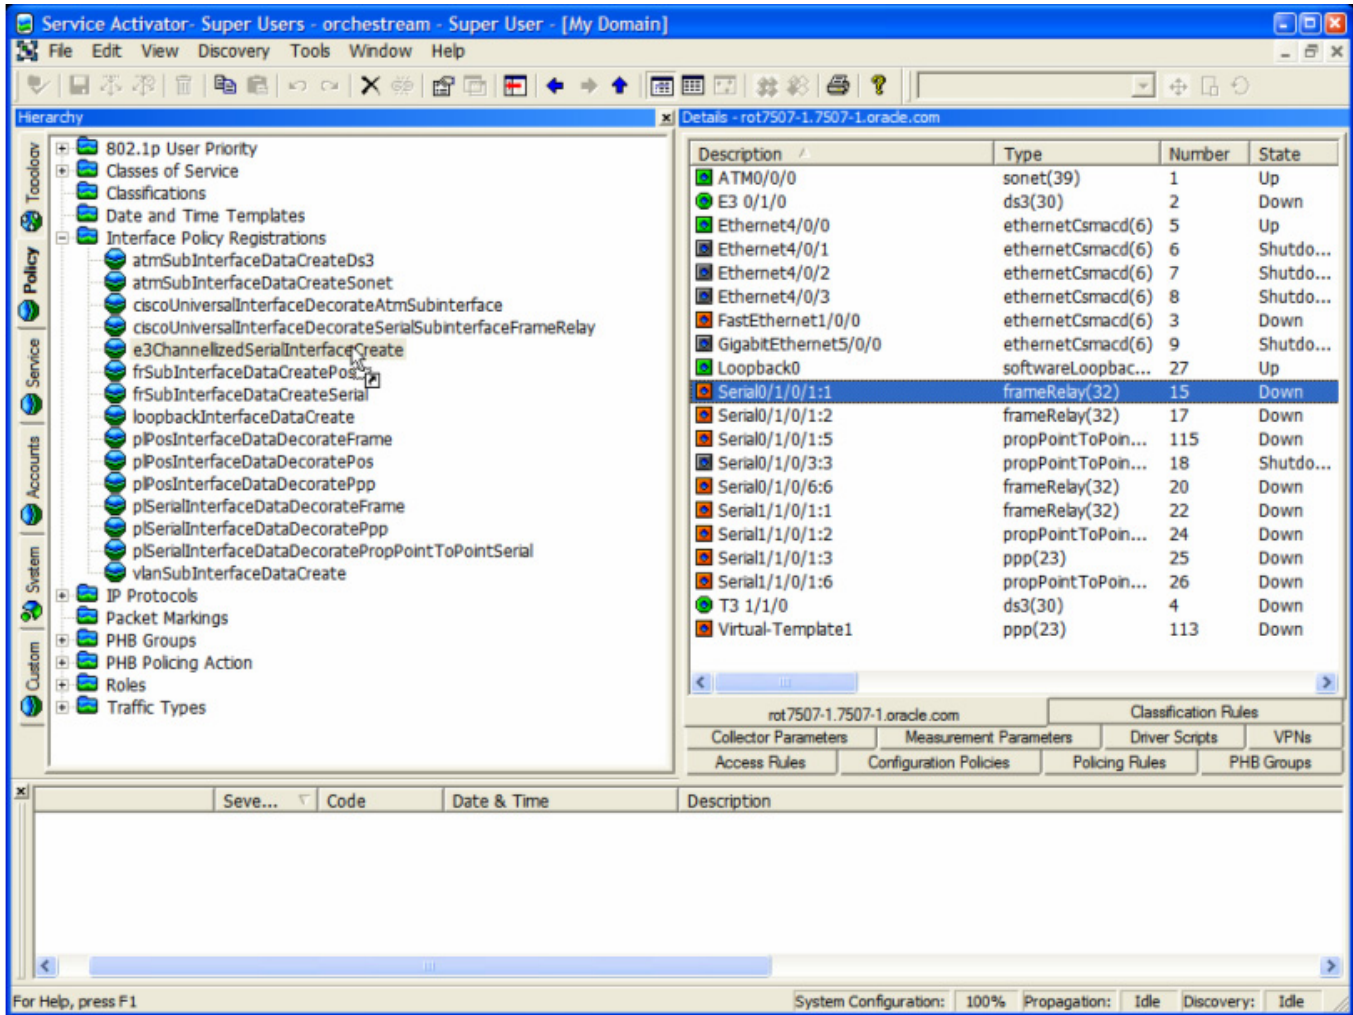
1. On the Policy tab right click **Interface Policy Registration** folder and select **Add Interface Registration**.
2. In the **Interface Policy Registration** dialog set the following attributes:
  - Name: `e3ChannellizedSerialInterfaceCreateFrameRelay`
  - Action: `Create`
  - Menu Text: `E3 Channellized Serial Interface Frame Relay`
  - Policy Type: `e3ChannellizedSerialInterface`
  - Context: `Device`
  - SNMP Enterprise ID: `9 (Cisco)`



3. Go to **Creation Template** tab and set the additional attributes
  - SNMP ifType: 32 (frameRelay)
  - Speed: 0
  - Interface Name Prefix: Serial

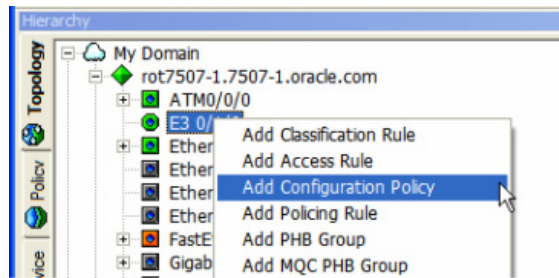


4. The Interface, Sub-Interface and VC Capabilities need to be defined. These are the capabilities that get assigned in the object model when the device level interface is created. These can be manually selected on the Interface Capabilities, Sub-Interface Capabilities and VC Capabilities tabs, or alternatively Capabilities for an existing interface can be used by dragging an existing and interface object onto the Interface Policy Registration Object.
5. Click **OK**, and drag a existing Serial interface onto the e3ChannellizedSerialInterfaceCreateFrameRelay interface registration policy, and Commit.

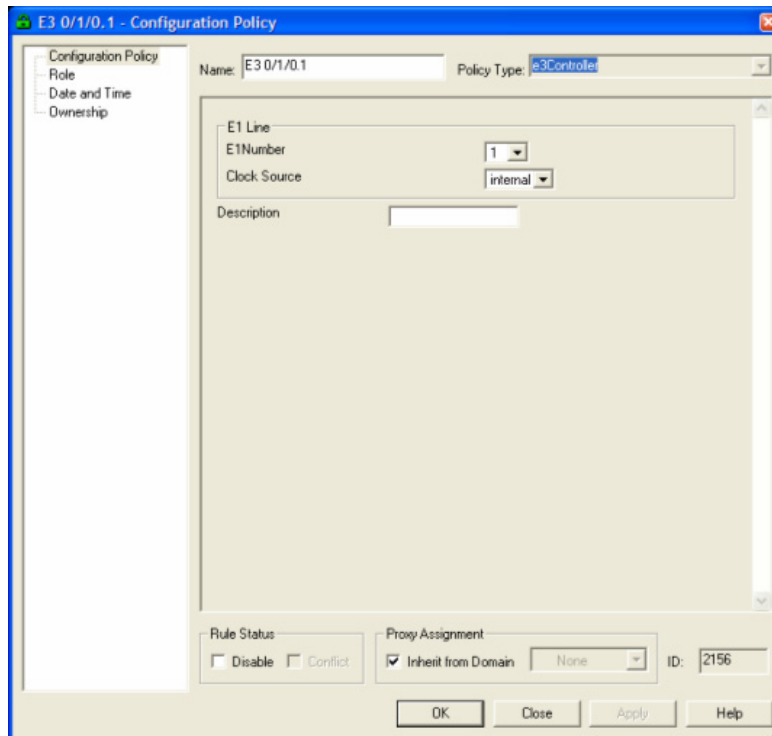


## Configuring the E3 Controller

1. Right click the E3 Controller and select **Add Configuration Policy** option.



2. The configuration policy dialog appears. Select Policy Type **3Controller** and provide a suitable descriptive name for the configuration policy. Set the policy attributes to configure the E1 line. Add a new configuration policy for each E1 line as described below
  - **Name:** E3 0/1/0.1
  - **E1 Number:** 1
  - **Clock Source:** internal



1. Set the configuration policy role, and click **OK**. Repeat these steps for each E1 Line to be configured and Commit.

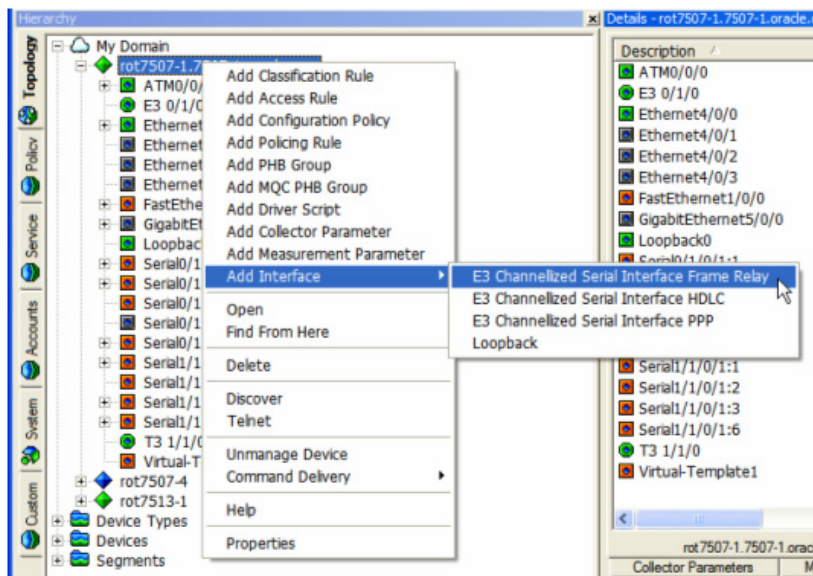
Name	Content Type
E3 0/1/0.1	e3Controller
E3 0/1/0.3	e3Controller
E3 0/1/0.6	e3Controller

In the above figure:

- controller E3 0/1/0
- e1 1 clock source internal
- e1 3 clock source internal
- e1 6 clock source internal

## Create the Channellized Serial Interface

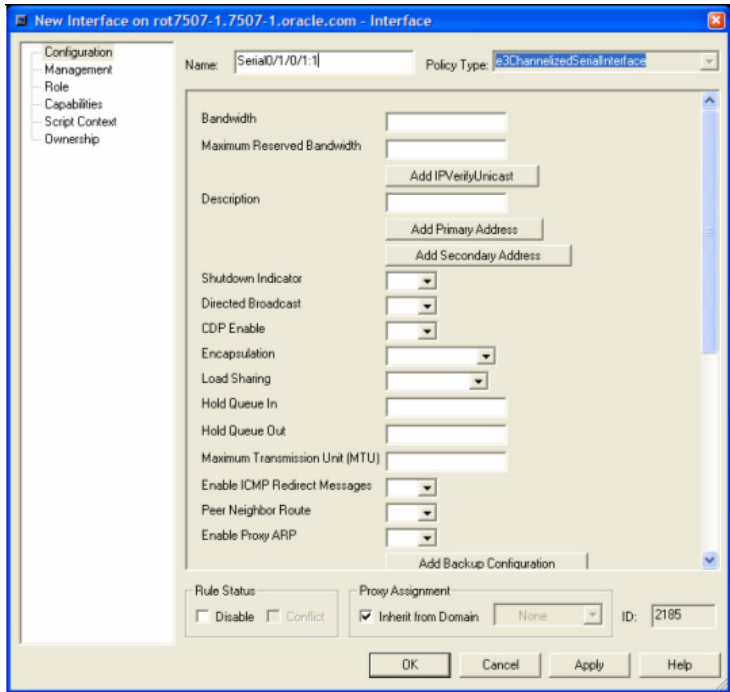
1. The e3ChannellizedSerialInterface interface creation policy is used to create a new Serial Interface on a device, based on the interface registration policy defined in [Interface Creation Policy Registration on page 158](#). To add the interface right-click the devices and select the **Add Interface>E3 Channellized Serial Interface Frame Relay**.



- The interface properties dialog will appear with the contents of the e3ChannellizedSerialInterface configuration policy. Set the interface data as appropriate. Note that the encapsulation field should be set to match the encapsulation applicable to the ifType defined in the interface registration policy to ensure the created interface capabilities are aligned.

propPointToPointSerial(22)	HDLC
ppp(23)	ppp
frameRelay(32)	Frame Relay Frame Relay IETF

- For the e3ChallenizedSerialInterface interface creation the Time Slots filed must be configured.
  - **Name:** Serial0/1/0/1.1
  - **Encapsulation:** Frame Relay
  - **Time Slots:** 1-5





4. Set the interface role, click **OK** and Commit.

IPSA will create the new interface on the device and configure allocate the time slots in the E3 controller.

```
2007-09-11 22:04:24|10.156.68.204|controller E3 0/1/0
2007-09-11 22:04:25|10.156.68.204|e1 1 channel-group 1 timeslots 1-5
2007-09-11 22:04:25|10.156.68.204|exit
2007-09-11 22:04:25|10.156.68.204|interface Serial0/1/0/1:1
2007-09-11 22:04:25|10.156.68.204|encapsulation frame-relay
2007-09-11 22:04:26|10.156.68.204|exit
```

## Removing channellized interface

Deleting the Interface from the IPSA object model will remove the Channellized interface from the device

```
2007-09-11 22:07:43|10.156.68.204|no interface Serial0/1/0/1:1.1
2007-09-11 22:07:43|10.156.68.204|interface Serial0/1/0/1:1
2007-09-11 22:07:43|10.156.68.204|no encapsulation
2007-09-11 22:07:43|10.156.68.204|exit
2007-09-11 22:07:43|10.156.68.204|controller E3 0/1/0
2007-09-11 22:07:44|10.156.68.204|no e1 1 channel-group 1
2007-09-11 22:07:44|10.156.68.204|exit
```

**Note:** any child sub-interfaces of the Channelized Interface will also automatically be removed, both from the object model and from the device.

## Hints and Tips

This section contains hints and tips for successful application of IPSA configuration policies.

### Device Level Configuration Policy Workaround

In order for a configuration policy to be successfully applied it must match a device and interface role and create a concrete at the interface, subinterface or VC object level. Some configuration policies are design to apply device level commands (i.e. not in the context of an interface) however to apply the configuration the configuration policy must be applied to an interface object in order to generate a concrete. Typically it is suggested that these device level configuration policies are attached to the Loopback interface.

The following is suggested as a usability workaround to allow Device level configuration policies to be declared at the Device level and automatically get concretes on the appropriate Loopback interface (or whatever interface is chosen).

1. Create a new interface role called **Device Policy**
2. Set the **Device Policy** Role on the selected interface (for example: Loopback0) on each device.
3. Add a Device level configuration policy to a Device
4. Set the Configuration Policy Interface Role to **Device Policy**

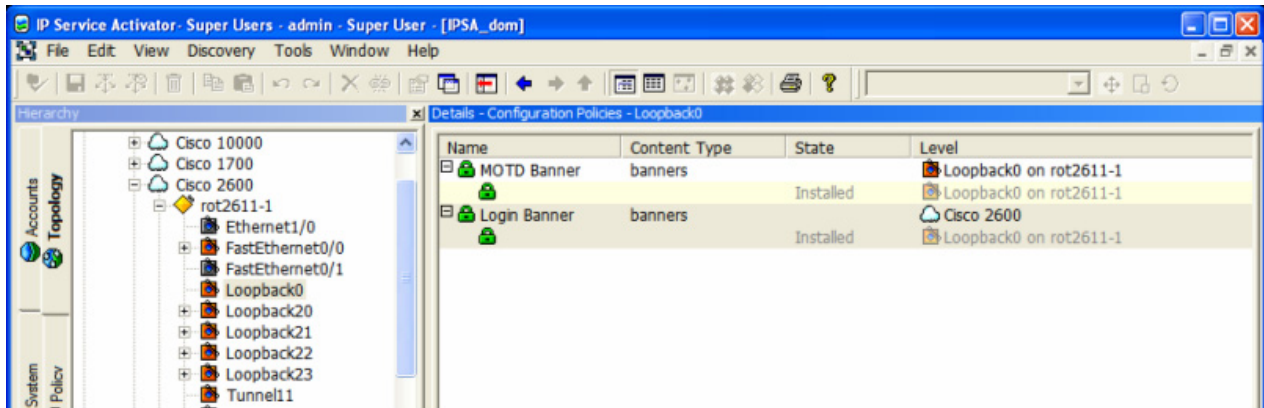
### Interface Creation and Decoration

An Interface Decoration Policy cannot be applied to an interface that has been created by a Creation Policy. To use interface management to take over management (and ownership) of the interface (i.e. to use the creation policy rather than the decoration policy), delete the sub-interface in the Object Model and re-create the same interface using the Interface Configuration Management.

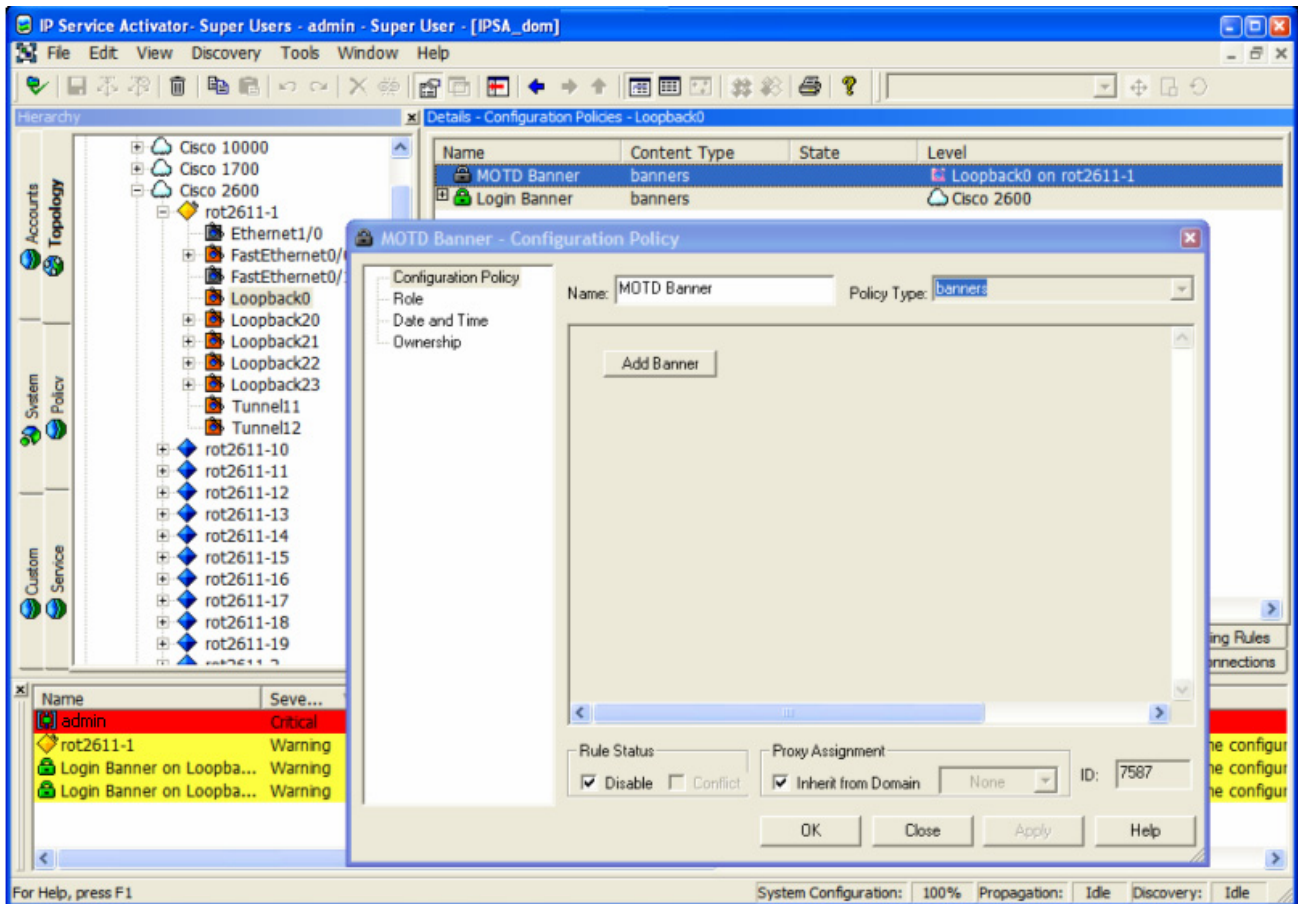
**Note:** As this deletes and recreates the object in the object model all service associations are lost and must be recreated and/or re-linked.

### Using Collection Based Configuration Policies

Some configuration policies such as the Static Route and banner policies allow for multiple definitions either within the same policies definition, or across multiple policy instances on the same device. This allows for creation of groups of definitions that may be defined at different levels in the policy inheritance hierarchy. For example, a general banner policy can be defined at the Network level with a common login and incoming banner definition that would apply to all devices, while a separate MOTD banner could be defined specifically at the device level.



Each collection based policies must contain at least one valid definition. If the last definition is removed , the policy must either be disabled or deleted from the device.



## Centralized Management of Configuration Policy GUI Extensions

By default, the configuration policy GUI extensions are manually installed on each IPSA client machine. This can be a cumbersome task when deploying and updating a large number of client installations. One approach to reduce the number of files updated on each client machine is to serve all of the GUI extension HTML files from a central location such as a shared drive or from a common http server.

This can be achieved by creating a master ConfigurationPolicy.cfg file and update the <policyurl> references for the required GUI Extensions to use a remote URL reference rather than a local file reference. E.g.

```
<GenericRule>
  <object>
    <policytypeReference ref="dlswDevice"/>
    <policytypeReference ref="dlswEthernetInterface"/>
    <policytypeReference ref="dlswTokenRingInterface"/>
  </object>
  <policytype name="dlswDevice">
    <policyurl>http://myserver/ipsa/dlswDevice.html</policyurl>
  </policytype>
  <policytype name="dlswEthernetInterface">
    <policyurl>http://myserver/ipsa/dlswEthernetInterface.html</
policyurl>
  </policytype>
  <policytype name="dlswTokenRingInterface">
    <policyurl>http://myserver/ipsa/dlswTokenRingInterface.html</
policyurl>
  </policytype>
</GenericRule>
```

The master ConfigurationPolicy.cfg must be deployed on each IPSA client installation. But updates to the configuration policy GUI extension html files for upgrades and patches only need to be updated in a single server location.



## Appendix C

# Oracle Communications Policy Services and IPSA Integration

Oracle Communications IPSA offers a standard licensed integration option for utilizing the power of Oracle Communications Policy Services in managing IP address assignments associated with Layer 3 MPLS VPN Sites. When provisioning VPN sites within IPSA, this integration allows the user to allocate IP addresses from Policy Services.

This integration facilitates the central allocation and tracking of IP address usage across VPNs to the customers using Oracle Communications Policy Services, a powerful application that provides comprehensive IP Address Management (IPAM), as well as DHCP, DNS and ENUM capabilities.

Within Policy Services IP address space is first divided into suitable IP subnets for allocation. IP subnets may be reserved on a per-customer or per-VPN basis. From within IPSA users can allocate IP subnets from Policy Services for MPLS IP VPN PE-CE point-to-point links.

Policy Services uses customer folders to organize IP address space, allowing overlapping private address space amongst customers if this is required. Each subnet has associated data for planning and tracking its usage including:

- Description
- Status (Free, Reserved, and Assigned)
- Custom Fields (Custom Fields 1 to 6) – specific names may be set within Policy Services

The integration uses the Status field to identify subnets available for allocation within IPSA. IPSA automatically updates this status when a subnet allocation takes place. Custom Field1 provides further grouping of IP subnets. For example, this field can be used to specify a Region or VPN that is allowed for this subnet. From an IPSA perspective, it is a read-only field. Custom Field2 is used by IPSA to associate additional Site information to the allocated IP subnet. This information includes the Site Name, PE router and PE interface. In addition, a “Remarks” field is supported in

IPSA that displays the IP subnet description field. This field can also be updated from the IPSA GUI when the subnet is allocated.

This integration allows you to select an IP subnet and mask for a VPN site from a list presented in the GUI. The IP subnets are associated with various IP address groups. From the IPSA GUI you can:

- Select the customer from which to allocate IP address space
- Select the Group from which to allocate address space
- Allocate an IP subnet from a specific Group.
- Choose which IP address within the subnet to associate with the PE interface
- De-allocate or release an IP subnet back into a Group.

## Integration Architecture

An intermediary server handles interaction between the Policy Service system and IPSA. The communications between IPSA and the intermediary server are performed using CORBA.

IP address Groups are configured using the Policy Services WebGUI client. The allocation/de-allocation of IP subnets is done via the Java CORBA server.

## Detailed Behaviour

When the IP Address Allocation GUI is launched within IPSA, the application retrieves the list of available customers from Policy Services. The application matches the IPSA customer with the returned list and pre-selects that customer. The user can override this pre-selection and re-select a customer from the drop-down menu. If there is no match, no pre-selection is done and **Please Select Customer** dialog is shown.

Once the customer has been selected, a list of available Groups is retrieved. The Group tag is stored in Customer Field1 in Policy Services. Once the user selects a Group, the available IP subnets associated with this Group are displayed. It is valid to have IP subnets not related to a Group. In this case, a blank group is listed amongst the available Groups. If there is only one group available, it is pre-selected.

The list of available IP subnets is displayed along with any supplemental data that has been associated with the IP subnet. This supplemental data can be entered using the Policy Services GUI during, for example, a pre-planning step for activating a Site. The Policy Services Description field is used to supply this additional data.

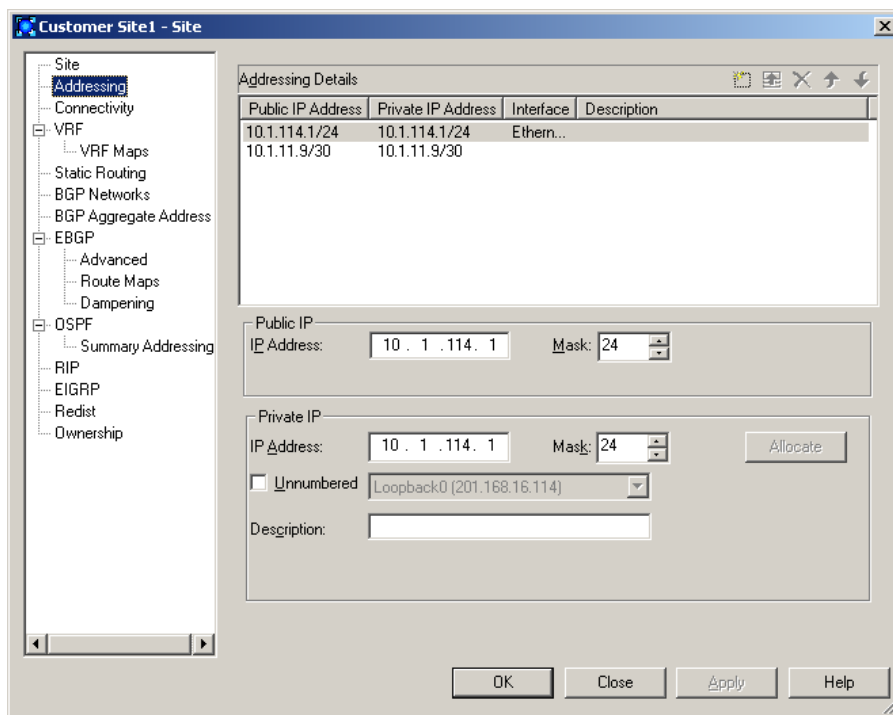
Once the IP subnet has been selected and the **OK** button is pressed, IPSA performs the following:



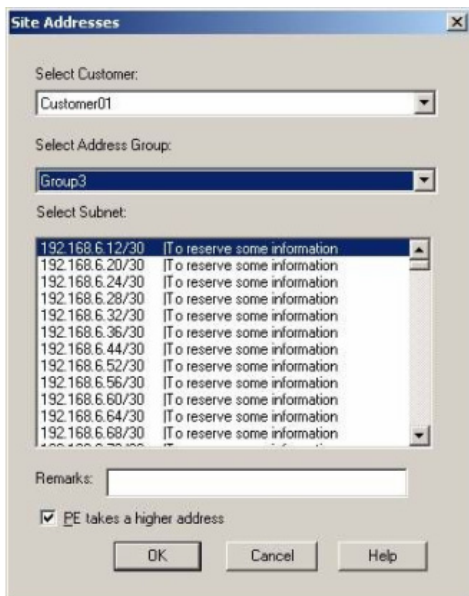
- Changes the status of the IP subnet to **Assigned**
- Writes the Remarks information into the Description field of the IP subnet
- Writes IPSA’s Site information into Custom Field2. This information includes the Site name, PE device and interface names. Any existing information in this field will be overwritten.

### Graphical User Interface Presentation

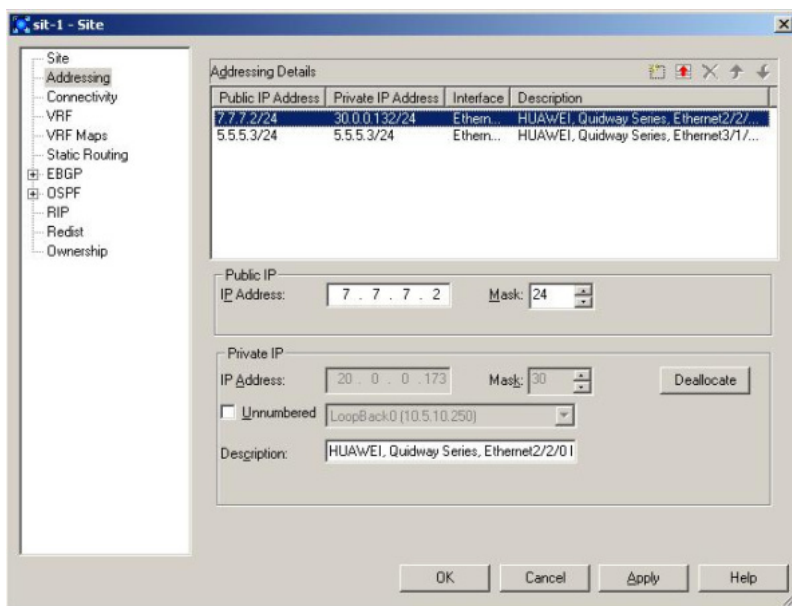
Allocation of an IP Subnet within IPSA is initiated from the Site object using the **Allocate** button in the **Addressing** property page



Click **Allocate** in on Site Addressing dialog to open the IP subnet selection dialog. Select a customer and an Address Group. The system queries Policy Services for available IP subnets in the selected Address Group and presents these subnets for selection.



From the list of available subnets, choose a specific IP subnet and add a remark in the **Remarks** field. Click **OK** to close the Address Group selection dialog. The **IP Address** and **Mask** fields in the IPSA are populated with the chosen IP address and mask and manual entry fields are disabled.



# Index

## Numerics

802.1Q, description of 77

## A

Access interfaces

in MPLS VPNs 19

in TLS 89

Access port

definition of 78

when configured in port and VLAN-based

TLS 86

Access rules, applying to CCCs 114

Advertisement Interval 28

Applying QoS to CCCs 113

AS\_PATH attribute

specify maximum occurrences of an ASN  
in 17

substitute ISP ASN in 17

ASN

and MPLS VPNs 5

core network ASN 18

override in AS\_PATH 17

specify maximum allowed in AS\_Path 17

Authentication

PE to CE 27

Autonomous systems and MPLS VPNs 5

## B

BGP and MPLS VPNs 5

## C

CCCs

applying QoS to 113

creating 108

deleting 112

implementing 110

Layer 2 switching 105

manual pre-configuration 106

MPLS tunneling 106

overview of 104

permitting/denying access to 114

queuing mechanisms 113

setting up 108

view implemented 112

CE router

linking to layer 2 site 96

linking to VPN site 24

CEF, prerequisite for MPLS 14

Circuit Cross Connects See CCCs

Community attribute, specifying standard or  
extended 17

Configuring CE routers, for MPLS VPNs 15

Configuring PE routers, for MPLS VPNs 14

Core interfaces, in MPLS VPNs 19

Core interfaces, in TLS 89

Creating

CCCs 108

customers 20

MPLS VPNs 46

sites 21

TLS 91

customer support ix

Customers, setting up 20

## D

dCEF, prerequisite for MPLS 14

Deleting

CCCs 112

documentation

downloading x

Service Activator xi

Domain settings, for MPLS VPNs 15

## E

eBGP

apply route dampening 28, 29

- parameters, configuring 26
  - See also Route redistribution
- Export maps, apply user-defined to VRF table 38
- Extended community attribute, use in BGP routes 17
- F**
- Fully-meshed MPLS VPNs, description of 3
- H**
- Hub and spoke MPLS VPNs, description of 3
- Hub site, specify in MPLS VPN 52
- I**
- IBGP and MPLS VPNs 5
- IBGP peer, specify maximum paths to 17
- IBGP peering
  - configure 17
  - pre-configured 17
- IBGP peers, use MD5 authentication between 17
- Implementing
  - CCCs 110
  - MPLS VPNs 55
  - TLS 99
- interface-less site 9
- interface-less VRF 9
- Interfaces
  - access, in MPLS VPNs 19
  - access, in TLS 89
  - core, in MPLS VPNs 19
  - core, in TLS 89
  - disabled in management VPNs 148
  - local, in MPLS VPNs 19
  - local, in TLS 89
- J**
- Juniper devices
  - access rules 114
  - CCCs 104
  - layer 2 VPNs 104
  - rate limiting 113
- L**
- Layer 2 Martini VPNs
  - Cisco devices 62
  - creating a VPN 71
  - description 59
  - devices and encapsulation types 62
  - Juniper devices 64
  - manual pre-configuration 68
  - modifying a VPN 72
  - overview of creation 65
  - provisioning endpoints (VC IDs) 70
  - technical description 60
- Layer 2 port, representation in user interface 79
- Layer 2 site
  - and CE access ports 85
  - and CE local ports 86
  - applying rate limiting to 97
  - associating physical component with 96
  - linking CE router to 96
  - linking to a TLS 96
  - port and VLAN-based 85
  - port-based 81
  - setting up 94
- Layer 2 switching CCCs 105
- Layer 2 VPNs 104
- Linking
  - CE devices to layer 2 sites 96
  - PE interfaces to layer 2 sites 96
  - port to layer 2 site 96
  - sites to MPLS VPNs 49
  - sites to VPNs 96
- Local ASN 27
- Local interfaces
  - in MPLS VPNs 19
  - in TLS 89
- Loopback addresses, and MPLS VPN set-up 14
- Loopback ID, define for MPLS VPN 17
- M**
- Management VPN
  - and QoS and SAA 12
  - function of 4
  - steps to set up 147
- Manual pre-configuration
  - CCCs 106
  - MPLS VPNs 13
  - TLS 89
- Map, creating a VPN 53
- Maximum paths, to iBGP peer 17
- MD5 authentication

- PE to PE 17
- Measurement, apply to MPLS VPNs 12
- MPLS and MPLS VPNs 2
- MPLS tunneling CCCs 106
- MPLS VPN topologies 3
  - fully meshed 3
  - hub and spoke 4
  - management 4
- MPLS VPNs
  - implementing 55
  - introduction to 2
- MPLS VPNs and QoS
  - applying QoS 12
  - order of setup tasks 46
- MPLS VPNs, advanced options
  - export map name 38
  - route distinguisher 37, 51
- MPLS VPNs, domain parameters
  - iBGP peering options 17
  - loopback ID, define 17
  - maximum paths to iBGP peer 17
  - send community attribute
    - BGP routes 17
  - setting up 15
- MPLS VPNs, management
  - QoS and measurement 12
- MPLS VPNs, manual pre-configuration
  - Export maps 14
- MPLS VPNs, measuring
  - applying measurement to 12
- MPLS VPNs, PE-CE connectivity 24
  - private PE to CE addresses 20
  - routing parameters 24
  - send community attribute
    - eBGP routes 28
  - Site of Origin 24
  - See also* eBGP, OSPF, RIP, Static routing
- MPLS VPNs, PE-CE routing parameters
  - setting up 24
- MPLS VPNs, route distinguishers
  - using VPN-wide 7
- MPLS VPNs, route targets
  - description of 47
- MPLS VPNs, routing protocols used in 5
- MPLS VPNs, setting up
  - associating components 22
  - creating 46

- creating a VPN map 53
- customers, setting up 20
- device discovery 19
- linking sites 49
- list sites in 54
- maps, creating 53
- overview of steps 13
- planning 3
- pre-configuring routers 13
- public and private addresses 19
- public and privates addresses 32
- setting up sites 21
- specify hub site 52
- specify topology 47
- using system-defined roles 19
- viewing implemented 57
- VRF tables, options for handling 6

MPLS VPNs, VRF tables

- manually pre-configured, handling 10
- options for handling 6
- re-use or reduction 8
- using VPN-wide table names 7

**N**

- Neighbour Description 27
- No Prepend 27

**O**

OSPF

- PE to CE routing protocol 25
- See also* Route redistribution

**P**

- PE interfaces, linking to sites 23
- PE-CE routing parameters
  - configure static routing 30
  - configuring eBGP 26
- PE-CE routing parameters, setting up 24
- Port
  - linking to layer 2 site 96
  - using in multiple port and VLAN-based TLSs 86
- Pre-configuring routers
  - MPLS VPNs 13
  - TLS 89
- Prefix filters
  - apply to VRF table 28
  - user-defined 14

- Prefix limit, apply to VRF table 28
- products
  - downloading x
- Propagating
  - CCC 110
  - TLS 101
  - VPNs 56
- Public and private addresses, in MPLS
  - VPNs 32
- Q**
- QoS policy
  - applying to CCCs 113
  - applying to MPLS VPNs 12
  - on MPLS VPNs, order of setup tasks 46
- Queuing mechanisms, applying to CCCs 113
- R**
- Rate limiting, applying to a layer 2 site 97
- RIP
  - select as PE to CE routing protocol 25
- Roles
  - assigning to devices in an MPLS VPN 19
  - Shadow role and VPN measurement 19
- Route dampening in eBGP 26
- Route dampening, applying in eBGP 28, 29
- Route distinguisher
  - override default 37, 51
  - supported formats 37, 51
- Route distinguishers
  - VPN-wide or site-specific 7
- Route redistribution, MPLS VPNs
  - filtering with route maps 14
- Route target numbers
  - setting 47
- Router-ID 10, 37
- Routes, specify maximum in VRF table 18
- S**
- SAA, applying to MPLS VPNs 12
- Securing TCP connections 17
  - PE to CE 27
- Service Application Point 9
- Setting up
  - CCCs 108
  - MPLS VPNs 46
  - MPLS VPNs, overview of steps 13
  - TLS 90
- site, interface-less 9
- Sites, layer 2
  - associating CE devices 96
  - associating PE interfaces 96
- Sites, MPLS VPN
  - associate interfaces with 22
  - associating CE devices 24
  - associating PE interfaces 23
  - fully meshed 3, 47
  - hub 4, 47
  - in VPNs 21
  - linking to an MPLS VPN 49
  - list MPLS VPN's sites 54
  - management 4
  - mesh. See fully meshed
  - mesh. See fully meshed
  - overriding CE ASN 17
  - route distinguisher 37, 51
  - route targets 47
  - setting up 21
  - spoke 4, 47
  - VRF Router-ID 10, 37
- Standard community attribute, use in BGP
  - routes 17
- Static routing
  - configure parameters 30
  - PE to CE routing 26, 30
- support
  - customer ix
- System-defined roles
  - in MPLS VPNs 19
  - in TLS 89
- T**
- Timers 28
- TLS
  - applying 99
  - creating 91
  - creating, overview of steps 80
  - device discovery 89
  - example setup 87
  - how frames are mapped to 77
  - implementing 99
  - linking layer 2 sites to 96
  - manual pre-configuration 89
  - overview of 76
  - planning 79
  - port and VLAN-based 82

- port-based 80
  - setting up 90
  - using port in multiple 86
  - using system-defined roles 89
  - viewing implemented 101
  - Topology of MPLS VPNs 3
  - Transparent LAN Service *See* TLS
  - Transport LSP, description of 77
  - Trunk port
    - when configured in port and VLAN-based TLS 86
  - Trunk port, definition of 78
- U**
- Unequal cost 37, 51
  - Update Source 27
- V**
- VC-LSP, description of 76
  - Viewing
    - implemented CCCs 112
    - implemented MPLS VPNs 57
    - implemented TLS 101
  - VLAN
    - definition applied to access port in TLS 85
  - VRF tables
    - advanced options for 36
    - apply domain VRF route limit 37, 51
    - default 37, 51
    - export policy 47
    - import policy 47
    - maintaining an interface VRF table 37, 51
    - manually pre-configured, handling 10
    - merging 37, 51
    - options for handling 6
    - prefix filters 28
    - prefix limit 28
    - re-use or reduction 8
    - route limit 37, 51
    - specify maximum number of routes in 18
    - user-defined 14, 37, 51
    - VPN-wide or site-specific names 7
  - VRF, interface-less 9
  - VRF-Aware IPsec 142

