

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

Juniper E-series Device Support Guide

Third 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	vii
About this document	vii
Before contacting Oracle Global Customer Support (GCS)	vii
Contacting Oracle Global Customer Support (GCS)	viii
Downloading products and documentation	viii
Downloading a media pack	ix
Service Activator publications	ix
Chapter 1 Summary of Features Supported	1
Supported devices	2
Supported operating system versions	2
Supported interface types	2
Virtual routers	3
Supported Service Activator features	4
Summary of features supported	4
MPLS VPN support	4
Chapter 2 Installation and Setup	7
Command-line parameters	8
Setting command-line parameters on Solaris	10
Setting command-line parameters on Windows	10
Chapter 3 Discovery and Configuration Issues	13
Device discovery issues	14
Address used to manage device	14
Discovery of large numbers of interfaces	15
How virtual routers are displayed	16
Obtaining device capabilities	16

Applying Service Activator configuration	17
The virtual device state	17
Check and force consistency	18
Dealing with manual configuration	19
Features and restrictions	20
Chapter 4 Manual Pre-configuration	21
Mandatory manual configuration for MPLS VPNs	22
PE routers	22
P routers	22
CE routers	23
Optional manual configuration for MPLS VPNs	23
Pre-defined VRF tables	23
Pre-defined export maps	24
Pre-defined prefix list filters	25
Manually pre-configured multi-AS VPNs	28
Chapter 5 Configuration of MPLS VPNs	31
Pre-configuration of routers	32
Domain-level parameters	32
Discovery and role assignment	34
Configuring VRF tables and route targets	34
VRF tables	34
Route distinguishers	36
RD number per VPN	36
VPN topology and route targets	37
VRF re-use/reduction	39
Co-existence with pre-defined VRF tables	39
Previously-defined export maps	40
Configuring PE-PE peering	41
Configuring iBGP	41
PE-PE community attributes	42
Co-existence with previously configured iBGP	42
MD5 authentication	42

PE-CE configuration using eBGP	42
Allow AS in	44
AS override	44
PE-CE community attributes	45
Authentication	45
Local preference	45
Site of origin	46
Route prefix limits and filters	46
eBGP load sharing	46
Route dampening	47
Route redistribution into eBGP	47
PE-CE configuration using RIP	48
Route redistribution into RIP	50
PE-CE configuration using static routing	51
Chapter 6 Troubleshooting	55
Communication problems	57
Useful E-series commands	58
General configuration	58
VPN configuration	58
Appendix A Sample Device Configuration	63
Sample configuration for a 3-spoke hub site	65
Index	69

Preface

About this document

The [Juniper E-series Device Support Guide](#) provides detailed technical information about the Juniper E-series (formerly Unisphere) device driver, including supported features, configuration requirements and detailed examples. It is intended for network managers and technical consultants responsible for implementing Oracle Communications Service Activator within a network using Juniper E-series routers.

The [Juniper E-series Device Support Guide](#) consists of the following chapters:

- [Chapter 1: Summary of Features Supported](#) explains the hardware, software and Service Activator features supported by the Juniper E-series device driver
- [Chapter 2: Installation and Setup](#) explains installation issues and details the device driver command line parameters.
- [Chapter 3: Discovery and Configuration Issues](#) explains the way in which Service Activator discovers and configures devices and ensures consistency.
- [Chapter 4: Manual Pre-configuration](#) details the prerequisites for running Service Activator including setting up routers for Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs).
- [Chapter 5: Configuration of MPLS VPNs](#) explains how MPLS-based VPNs are implemented, including details of commands configured.
- [Chapter 6: Troubleshooting](#) gives hints and tips for diagnosing and fixing problems.
- [Appendix A: Sample Device Configuration](#) includes example configurations of routers.

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. In order to obtain the broadest access to skilled technical support, Oracle advises you to log 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

Summary of Features Supported

This chapter outlines Service Activator’s support for Juniper E-series devices. It includes the following:

- Details of the Juniper E-series hardware and software supported by this release of Service Activator, including information on operating system versions, interface types and the E-series virtual router concept
- Details of the Service Activator features that are supported by the Juniper E-series device driver

Supported Juniper E-series hardware and software

The Juniper E-series device driver is effectively capable of configuring any E-series device. However, the exact capabilities that can be supported depend on the device model, the operating system that it is running and the interface.

If you are using, or wish to use, different hardware or software from that defined here, please contact Technical Support for more information.

Supported devices

The E-series comprises the following models:

- ERX-705
- ERX-710
- ERX-1410
- ERX-1440

Supported operating system versions

Oracle's policy is to guarantee support for only those operating system versions against which Service Activator has been thoroughly tested. Refer to the *Release Notes* for your Service Activator software release for support information for the specific releases of the JUNOS operating system.

For up-to-date information on functionality supported in these operating system versions, please see the E-series page on the Juniper website (www.juniper.net).

Supported interface types

The Juniper E-series device driver supports a wide range of interface types, as summarized below:

- Gigabit Ethernet
- VLANs
- Channelized T3
- Unchannelized T3/E3
- Channelized T1
- Channelized E1
- Dual Port OC3
- Ethernet

- ATM
- Frame Relay
- Multilink Frame Relay
- PPP
- PPP over Ethernet
- Packet over SONET
- Bridged IP
- Cisco HDLC

For full details of which interfaces are supported on specific devices, please consult the router documentation or the E-series page on the Juniper website (www.juniper.net).

Virtual routers

Juniper E-series hardware supports the concept of virtual routers, which allows a single physical router to behave as several separate routers.

Virtual routers are totally separate entities — they behave as routers in their own right and run their own BGP and OSPF routing instances.

Every E-series physical device has one default virtual router. When you first log into an E-series device you are accessing the default virtual router. The default virtual router behaves exactly as other virtual routers except that it cannot have child VRF virtual routers.

A virtual router may have one or more child VRF virtual routers, which represent VRF tables. VRF virtual routers are always created as children of virtual routers but not of default virtual routers. VRF virtual routers do not have their own BGP and OSPF routing instances, and they cannot have child VRF virtual routers themselves.

For details of how virtual routers are discovered and displayed on the Service Activator user interface, see [Discovery of large numbers of interfaces on page 15](#).

Virtual routers should not be confused with virtual devices which are created within the user interface to represent devices that Service Activator has not discovered.

Supported Service Activator features

Summary of features supported

- Service Activator fully supports the E-series virtual router concept. Virtual routers can be discovered, managed and configured.
- Service Activator fully supports the provisioning of MPLS VPNs on Juniper E-series devices.

MPLS VPN support

Service Activator feature		JUNOSe 3.4.1	JUNOSe 4.1.0
VRF table	User-defined VRF table name	✓	✓
	VRF re-use	✓	✓
	User-defined RD numbers	✓	✓
	User-defined RT numbers	✓	✓
	VRF route limit (max routes)	✓	✓
	Co-existence with pre-defined VRFs	✓	✓
	Pre-defined export maps	✓	✓
PE-PE peering (iBGP)	iBGP peering optional	✓	✓
	Maximum paths	x	x
	Extended/standard community attributes	✓	✓
	PE-PE MD5 authentication	✓	✓
PE to CE connectivity	eBGP	✓	✓
	OSPF	x	x
	RIP	x	x
	Static routing	✓	✓

Service Activator feature		JUNOSe 3.4.1	JUNOSe 4.1.0
eBGP configuration	AS override	✓	✓
	Allow AS in	✓	✓
	Extended/standard community attributes	✓	✓
	Local preference	✓	✓
	Prefix filters	✓	✓
	PE-CE authentication	✓	✓
	Prefix limit	✓	✓
	Site of origin	✓	✓
	Multi-path load sharing	✓	✓
	Route dampening	✓	✓
	Route redistribution into eBGP	✓	✓
Static configuration	Global routes	✓	✓

Chapter 2

Installation and Setup

This chapter explains how to set up and run the Juniper E-series driver. It includes the following:

- Installation issues
- Details of the command-line parameters that can be used when setting up the Juniper E-series device driver

Installation issues

The Juniper E-series device driver is a separate executable component. Note that the executable is currently called **unisphere_device_driver.exe**

The Juniper E-series device driver is always installed when you select the Proxy Agent install. For more information see the *Setup Guide*.

Command-line parameters

The following table summarizes the command-line parameters recognized by the Juniper E-series component.

Parameter	Description
-ComponentName < <i>name</i> >	Specifies the name of the Juniper E-series device driver component as displayed in the user interface.
-ComponentLocation < <i>hostname</i> >	Specifies the hostname on which the Juniper E-series component is installed.
-TelnetPort	Set the port the driver uses to communicate with the device. The default port is 23.
-NumRetries < <i>n</i> >	Number of times to retry a socket connection attempt, where <i>n</i> is an integer. Default is 0.
-ConnectTimeout < <i>n</i> >	Socket connection timeout, where <i>n</i> is a integer specifying number of seconds. Default is 30 seconds.
-ReadTimeout < <i>n</i> >	Socket read timeout in seconds, where <i>n</i> is a integer specifying number of seconds. Default is 30 seconds.
-WriteTimeout < <i>n</i> >	Socket write timeout in seconds, where <i>n</i> is a integer specifying number of seconds. Default is 30 seconds.
-NoCommandDelivery	Commands are logged but not sent to any devices.

Parameter	Description
-ForceVpnRollback	Tests VPN rollback. Causes the driver to issue all VPN commands, roll them all back and then re-issue them
-DisableVpnPreservation	Can be used when Service Activator is used to apply all VPN configuration. Any manual VPN configuration detected on the device will be removed.
-ParserCmdTimeout <time>	Specifies the configuration timeout period, in seconds. The driver retrieves the router's configuration at this interval. This defaults to 10 minutes. It is not recommended that you do not change this setting unless advised to do so by Technical Support.
-ExcludeCategories <"string">	Allows the exclusion of sections of configuration that are not to be parsed by the device driver. See Discovery of large numbers of interfaces on page 15 .
-OnlyMatchInterfaceType <regex>	Allows the user to limit the interface types that are to be parsed by the device driver by means of a regular expression. See Discovery of large numbers of interfaces on page 15 .
-OnlyMatchInterfaceSlot <regex>	Allows the user to limit the interface slots that are to be parsed by the device driver by means of a regular expression. See Discovery of large numbers of interfaces on page 15 .
-UseMechanismFile	Forces the driver to use capabilities files rather than the statically compiled versions.
-MechanismDirectory <path>	Specifies the location of the capabilities files.
-MechFailOnError	When this flag is enabled, if an error is encountered in the capabilities files a failure is reported, otherwise the driver will ignore the error and continue processing.
-AlwaysPass <regex>	Regular expression containing patterns for all the extra return messages which should pass.

Parameter	Description
-AlwaysFail <regex>	Regular expression containing patterns for all the extra return messages which should fail.
-NoExpressCommandGeneration	Specifies the command generator mode. By default, if the option is not specified, the mode is express. If specified, the mode is normal. The command line option is linked to a component parameter with the same name, so it can be changed dynamically. The parameter's value is taken into account when a device (virtual router) is being managed. Changing the value of the parameter has no impact on already managed virtual routers. This allows for having a mixture of virtual routers, some managed in express mode and some in normal mode.

Note that there are also command-line parameters that control debugging logs for all Service Activator components. These are described in full in the *Administrator's Guide*.

The device driver must be restarted for any changes to these command-line parameters to take effect.

Setting command-line parameters on Solaris

Command-line options are specified in the **cman.cfg** file which is in the **/opt/OracleCommunications/Service Activator/Config** directory.

Using a text editor such as vi, edit the **unisphere** entry in the **cman.cfg** file with the relevant option.

Setting command-line parameters on Windows

Command-line options are specified in the registry entry.

1. Start regedit.
2. Browse to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\Orchestream\DPE\2.0\ Component Manager
3. In the right-hand pane, right-click on **unisphere** and select **Modify** from the pop-up menu.

The Edit String dialog box opens.

4. Edit the **Value data** field to set the relevant option.

Chapter 3

Discovery and Configuration Issues

This chapter explains how the Juniper E-series device driver communicates with and configures devices. It includes the following sections:

- How the driver communicates and authenticates with devices
- The discovery and representation of Juniper E-series devices, with particular reference to virtual devices
- How the device driver models and applies device configuration and ensures consistency of configuration
- Features and restrictions

Communication and authentication

The Juniper E-series device driver accesses devices via the command-line interface (CLI). Access is authenticated via local passwords. You must ensure that the authentication methods are correctly set up for all E-series devices in your network. You can set the authentication method on the **Security** property page in the **Discovery** dialog box to ensure that it applies to all devices, or set it for individual devices.

Note that the Juniper E-series device driver requires write access to devices in order to return capabilities. Ensure that device security parameters are set up correctly to obtain capabilities.

Device discovery issues

Address used to manage device

To ensure that Service Activator is able to apply configuration to E-series virtual routers, the address used to manage the device must not be changed from that used for discovery. This is because the IP address used must be the one valid for the default virtual router. If an alternative address is used, the device can be discovered but the driver cannot communicate with it.

You must set the correct option before discovery.

To set the address selection option before discovery

1. Select **Options** from the **Tools** menu. The **Options** dialog box opens.
2. Select the **Discovery** property page.
3. For both the **CE Device IP Address Selection** and **Non CE Device IP Address Selection**, select **Do not change**.

To confirm the address selection option after discovery

1. For each virtual router discovered, open its Properties dialog box by right-clicking on the virtual router on the **Topology** tab and selecting **Properties**.
The **Properties** dialog box appears.
2. On the **Device** page, check that the **IP Address** is blank. If it is not blank, select the blank option in the drop-down menu.
3. Click **OK** to close the **Properties** dialog box.

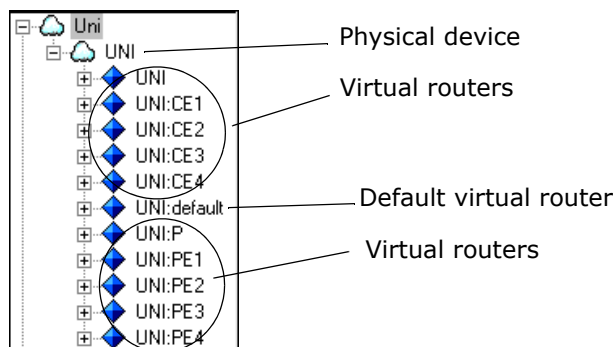
Discovery of large numbers of interfaces

When discovering large numbers of interfaces on virtual devices, it is advisable to limit the configuration parsed by the device driver. Three command-line options are available:

Parameter	Description
-ExcludeCategories <"string">	When virtual devices have large numbers of interfaces, this allows the exclusion of sections of configuration that are not to be parsed by the device driver. The string must be the exact command to add to show config commands to restrict the configuration. For example, the following command-line option ensures that ATM interfaces are not parsed: -ExcludeCategories "exclude-category interface atm"
-OnlyMatchInterfaceType <regex>	When virtual devices have large numbers of interfaces, this allows the user to limit the interface types that are to be parsed by the device driver by means of a regular expression.
-OnlyMatchInterfaceSlot <regex>	When virtual devices have large numbers of interfaces, allows the user to limit the interface slots that are to be parsed by the device driver by means of a regular expression. When using this option, it is essential that the regular expression matches all parent interfaces of the sub-interface required. For example, the following command matches slot 5/0.3 and 5/0.4: -OnlyMatchInterfaceSlot "^5/0(\.[34])? \$"

How virtual routers are displayed

E-series virtual routers are displayed in a similar way to real devices. When a device containing a virtual router is discovered, Service Activator creates a network cloud to 'hold' the routers. Beneath this network cloud, Service Activator creates device objects that represent both the physical device and its virtual routers. The physical and the virtual routers exist at the same hierarchy level, as shown below.



Virtual routers are identified in the user interface by the following name format:

physical-router-name:virtual-router-name

Default virtual routers are identified in the user interface by:

physical-router-name:default

VRF virtual routers are identified in the configuration by the following name format:

virtual-router virtual-router-name:Orch_RD-number.

Note that VRF virtual routers are not displayed in the hierarchy pane of the user interface.

Because virtual routers can never be independent of the physical router, you should never move virtual routers on the user interface, or drag and drop them between networks.

Obtaining device capabilities

The features supported on each device are dependent on the hardware and software combination and are indicated by a list of capabilities returned from each device.

You should always confirm capabilities before configuring network devices.

For E-series devices, specific capabilities are returned for each virtual router. For the physical router, capabilities will indicate that nothing is supported. The physical router should not be managed as no configuration can be applied directly to it.

Capabilities are obtained at two levels:

- At device level, based on the device type and OS version, the capabilities specify the device's support for measurement options, Transparent LAN Services and VLANs. As these are not supported on E-series devices the device-level capabilities are always blank.
- At interface, sub-interface and VC endpoint level, the capabilities indicate the MPLS features that can be configured on that interface.

The Juniper E-series device driver requires write access to devices in order to return capabilities. Therefore you need to ensure device security parameters are set up correctly first. The device driver accesses devices via the command-line interface, authenticating access as an anonymous user via local passwords.

The device driver attempts to obtain the capabilities automatically at the end of the discovery process. If this fails – for example, because the device's security settings are incorrect – you can initiate the process to fetch capabilities manually.

In addition, if the device capabilities have changed – for example, as a result of an operating system upgrade – you can reset the capabilities and re-manage the device.

For more details, see the *Network Discovery and Basic Setup* guide.

Do not use the loopback address for discovery of Juniper ERX VRs. This will cause the capabilities fetch to fail.

To avoid this, access the **Options** dialog box by selecting **Tools > Options** from the main menu. Select the **Discovery** property page and select **Do not change** for both the **CE Device IP Address Selection** and the **Non CE Device IP Selection**.

Applying Service Activator configuration

The virtual device state

The Juniper E-series device driver maintains a virtual version of the device state, consisting of all the configuration that has been defined.

Although the virtual state is an abstract representation of the real state, there is a direct mapping between the two. A real state can always be generated from a virtual state.

The device driver must ensure that the real state is a representation of the virtual state. It does this by extracting the real state of the device, comparing this with the virtual state and generating a set of commands that will ensure the real state matches the virtual state.

A sequence of steps is run every time the virtual state changes, that is, whenever a transaction is committed and changes are propagated.

- The device driver extracts the real configuration from the device, and compares this with the virtual device state.
- If configuration exists in the virtual state but not in the real state, it is installed.
- If configuration exists in both, the device driver checks that both are the same. If they are, no changes are made. If they aren't the same, the real state is updated.
- If configuration exists in the real state but not in the virtual state, the device driver's action depends on the setting of the **Manual Config Warning** flag and the VPN preservation setting. See [Dealing with manual configuration on page 19](#).

Check and force consistency

A check and force consistency process ensures that the configuration of each device always matches the virtual device state, even if it has failed and configuration is lost.

At regular intervals, the proxy agent polls the devices controlled by the driver. If it finds that the device was down (if it has re-booted since the last time it was checked), the proxy agent tells the driver to check the consistency of the device configuration.

If the driver finds that the real state is not as expected, it immediately issues commands to the device to bring the real state into line. If the driver finds that the real state matches the virtual state, it does nothing.

In this way the system makes sure that if any changes are made to devices, or if devices fail and re-boot, that the configuration is automatically reset. This does not apply if connectivity is lost but there is no device reload.

Dealing with manual configuration

You can set a **Manual Config Warning** flag that determines how the Juniper E-series device driver deals with manual configuration discovered on routers it is managing. The following options are available:

- **Warn and delete:** A warning message (3203 - Manual configuration detected) is output. Any unexpected manual configuration discovered by Service Activator is removed automatically, but manually configured VRF options are by default maintained.
- **Fail and don't delete:** A critical fault (3494 - Manual configuration detected - configuration not applied) is output, and the device status is set to Intervention Required. The manual configuration is not removed from the device, but the commands that would have been removed are saved in the device log file in the AuditTrails directory. No new configuration is applied.
- **Delete:** If manual configuration is discovered, no warning or fault is raised. Any unexpected manual configuration discovered by Service Activator is removed automatically, but manually configured VRF options are by default maintained. This is the default setting.

The default driver behavior can be set on the **Domain** dialog box. It can be overridden for specific devices if required by a setting on the **Device** dialog box.

By default, the removal of manual configuration does not apply to VRF tables on the device, as it is possible for Service Activator's MPLS VPN configuration to co-exist with pre-defined VRF tables or manually-applied commands within the VRF address-families.

However, you can force Service Activator to automatically remove all pre-defined VRF configuration that may exist on devices, by using the device driver command-line parameter:

```
-DisableVpnPreservation
```

Note the following:

- Manual configuration is only detected when the driver sends a configuration update to the device. If the virtual configuration does not change and the device never goes down, the driver will not update the device configuration and manual configuration will not be detected.
- The device driver only detects changes to configuration relating to those features supported by Service Activator.
- It is strongly recommended that manual configuration of policy features controlled by Service Activator is not permitted.

Features and restrictions

Note the following:

- The device driver only ever changes the configuration of those interfaces controlled by Service Activator.
- The driver removes existing configuration only if essential and at the last possible moment.
- If configuration exists on an interface/device that conflicts with the virtual configuration, but which is essential for routing/device operation, it will not be removed or changed.
- Configuration installed by the device driver is identified by objects being given specific names.

Chapter 4

Manual Pre-configuration

This chapter specifies the prerequisites for running Service Activator, and details the commands that you need to apply to ensure that Juniper E-series devices are properly set up and ready to be managed by the Juniper E-series device driver. It includes the following:

- Configuring SNMP for Service Activator discovery
- Pre-requisite manual pre-configuration for MPLS VPNs
- How Service Activator can co-exist with manually configured MPLS VPN features

Configuring SNMP

SNMP must be enabled on all E-series routers for the Service Activator discovery process to work. Ensure the following line is included in the physical device configuration:

```
snmp-server community public RO
```

Note that Service Activator's network discovery process uses a default community of `public`; you will need to amend the appropriate SNMP parameter in the **Discovery** dialog if you set a different read community on the physical devices.

Mandatory manual configuration for MPLS VPNs

Before using Service Activator to configure VPNs, some manual configuration of E-series physical devices is required.

In versions of the E-series operating system earlier than 3.2 you have to manually set up MPLS tunnels. In later versions, there is a configuration option to ensure that this is done automatically. See E-series documentation for details.

PE routers

On all PE (gateway) routers in the core VPN, the following manual configuration is required:

- IP addresses must be correctly assigned.
- A loopback interface must be set up.
- MPLS must be enabled on all appropriate interfaces by running the `mpls` and the `mpls ldp` commands.
- An IGP must be implemented in order to distribute IP routes.

P routers

On all P routers in the core VPN, the following manual configuration is required.

- IP addresses have to be correctly assigned.
- Loopback interfaces may be set up.
- MPLS must be enabled on all appropriate interfaces by running the `mpls` and `mpls ldp` command.

CE routers

The CE (access) routers at customer sites are not configured to control routing 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 the following are set up:

- BGP or static routing must be configured in order to advertise reachability information between the CE and the PE.
- It is recommended that a loopback interface is set up on each CE router.

Optional manual configuration for MPLS VPNs

It is possible to set up Service Activator to co-exist with manually configured VRF tables and export maps. This section explains how these features must be set up in order to co-exist correctly with configuration generated by Service Activator.

Pre-defined VRF tables

You can manually configure VRF tables on a PE router. When a pre-defined VRF table exists on a device, Service Activator can treat it in three different ways:

- Service Activator has no control of the VRF table or its contents
- Service Activator has control of the VRF table and preserves its contents
- Service Activator has control of the VRF table and removes its contents

You specify the amount of control Service Activator has over a VRF table by setting certain site-specific values (on the **VRF** page of the **Site** dialog box). For more information, see [Co-existence with pre-defined VRF tables on page 39](#).

Restrictions

VRF table name

The name of a user-defined VRF table must be unique on the device. It may consist of up to a maximum of 30 alphanumeric and underscore characters.

Adding route targets and external features

You can manually add route targets and parameters that cannot be configured by Service Activator to a pre-defined VRF table that is controlled by Service Activator. However, these parameters are preserved only until Service Activator either deletes the VRF table or merges it into another one. This normally occurs if you change the property settings of the relevant site, in which case the manually added route targets and parameters are no longer required.

Device driver restart

If the device driver has to be restarted, it must be reminded which VRF tables were created by Service Activator and which are pre-defined. This is implemented by re-propagating the configuration to the device driver. The configuration that you propagate must be the same as that propagated to the device since the last successful transaction before the device driver failed. The device driver responds to the update by classifying VRF tables on the device whose names match those in the user interface as Service Activator controlled, and the remaining VRF tables as uncontrolled. Any route targets or unknown parameters in a VRF table are preserved. If the configuration has changed since the last successful transaction, Service Activator cannot correctly identify all VRF tables on the device.

Pre-defined export maps

You can manually pre-define an export map on a PE router and assign it to a VRF table on the **VRF** page of the **Site** dialog box. The export map only allows those routes in the VRF table whose route prefixes match those specified in the export map to be advertised to other PE routers. The exported routes are tagged with an RT value specified by the export map.

Configuring an export map

The following commands provide an example of configuring an export map.

```
access-list 1 permit 128.1.1.1
```

Defines access list 1 which accepts routes having IP address 128.1.1.1

```
access-list 2 permit any
```

Defines access list 2 which accepts any routes

```
route-map export-map-name permit sequence-number
```

```
match ip address 1
```

```
set extcommunity rt 100:94
```

Export map *export-map-name* attaches route target 100:94 to routes specified in access list 1. The *sequence-number* identifies the order in which the route-map is implemented.

```
route-map export-map-name permit sequence-number
```

```
match ip address 2
```

```
set extcommunity rt 100:26
```

Export map *export-map-name* attaches route target 100:26 to routes specified in access list 2. The *sequence-number* must be a higher value than the preceding *sequence-number*.

If an export map is used by a management VPN, the spoke sites are not required to export route targets. To prevent management spoke sites exporting route targets, set **Spoke** to **None** for the spoke site's export policy route target on the **MPLS** page in the **VPN** dialog box, or alternatively, use the pre-defined export map configuration shown in the following example:

```
export map 'ExpMapCust#1'  
route-target export 1:1111  
route-target export 1:1394  
route-target export 1:1614
```

where ExpMapCust#1 is a pre-defined export map used by both management and customer VPN sites; 1:1111 is the route target of the management hub site, 1:1394 and 1:1614 are the route targets of the customer sites in the VRF table of each spoke site.

```
ip access-list extended ExpMap_Mng  
deny ip 192.168.65.0.0.0.0.255 any  
deny ip 20.20.20.0.0.0.0.255 any  
permit ip any any
```

where deny ip 192.168.65.0.0.0.0.255 any rejects matching routes to the management hub site, deny ip 20.20.20.0.0.0.0.255 any rejects routes to the customer LAN, permit ip any any accepts all other routes

```
route-map ExpMapCust#1 permit 10  
match ip address ExpMap_Mng  
set extcommunity rt 1:1394 1:1614
```

Export map ExpMapCust#1 attaches route targets 1:1394 and 1:1614 to routes permitted by access list extended ExpMap_Mng. Note that the management hub site route target 1:1111 is not attached to these routes.

Pre-defined prefix list filters

The number of routes that are received from, or sent to, a CE router can be selectively reduced using a manually pre-defined prefix list installed on the neighboring PE router. Routes whose prefixes match those in the prefix list will either be allowed or rejected by the PE router depending on their designation in the prefix list. You need to specify in the user interface that the prefix list is required to only filter routes that are either incoming (CE to PE) or outgoing (PE to CE).

A pre-defined prefix list can be used instead of an access list for configuring a pre-defined export map described in [Pre-defined export maps on page 24](#).

Creating a prefix list

You configure a prefix list in router configuration mode using the commands described below. You apply a pre-configured prefix list filter to a site by entering the name of the prefix list in the **Prefix filters In** or **Out** fields in the **Site** dialog box, **EBGP Adv.** page.

Arguments shown within brackets [] are optional.

This command string adds a single entry to a prefix list.

```
ip prefix-list list-name [seq sequence-value] deny/permit prefix/  
prefix-length [ge ge-value] [le le-value]
```

Sequence values are automatically generated by default. You only need to specify a sequence value if the automatic generation of sequence values is disabled. For more information, see [Sequence values](#).

ge-value and *le-value* specify a prefix length range, where:
prefix length < *ge-value* <= *le-value* <= 32.

Sequence values

Sequence values are generated automatically by default, but generation can be disabled using:

```
no ip prefix-list sequence-number
```

Sequence values are, by default, automatically generated in increments of 5, so that the first list entry has a value of 5 and the next entry has a value of 10, and so on.

If the automatic generation of sequence values is disabled, it can be re-enabled by:

```
ip prefix-list sequence-number
```

If generation of sequence values is disabled, you must specify a sequence value for each subsequent entry using the `seq` argument. You can specify any sequence value for subsequent entries. Note that incrementing each entry by 1 will prevent additional entries to be added between two existing entries at a later date. If you specify a sequence value for an entry but do not specify the value for the next entry, the value of the next entry is incremented by 5. For example, specifying a sequence value of 9 will make subsequent entries have sequence values of 14, 19, 24, and so on.

How prefix list filtering works

If a route prefix received by the PE matches a prefix in the prefix list, that prefix will either be accepted or rejected depending on whether the entry is designated as **permit** or **deny**. The following conditions also apply:

- A prefix is denied if it cannot be matched with any prefixes in the prefix list

- If a prefix matches several prefixes in the prefix list, the prefix with the lowest sequence value is used

Prefix list filter configuration command examples

The following examples use a prefix list called filter1.

Deny prefix

Deny default route 0.0.0.0/0:

```
ip prefix-list filter1 deny 0.0.0.0/0
```

Deny routes with prefixes 196.0.0.0/8 and prefix lengths greater than 25 up to 32 in network 196/8:

```
ip prefix-list filter1 deny 196.0.0.0/8 ge 25
```

Deny all routes in Class A network 22/8 by specifying prefix lengths from /8 to /32:

```
ip prefix-list filter1 deny 22.0.0.0/8 le 32
```

Deny routes with prefixes 100.70.1/ with prefix lengths from /24 to /25:

```
ip prefix-list filter1 deny 100.70.1.0/24 ge 25
```

Permit prefix

Permit route 36.0.0.0/8:

```
ip prefix-list filter1 36.0.0.0/8
```

Permit routes with prefix lengths of 8 to 24; make list entry sequence value 5:

```
ip prefix-list filter1 seq 5 permit 0.0.0.0/0 ge 8 le 24
```

Permit routes with prefixes 129.0.0.0 having prefix lengths of 8 to 24 in network 129/8:

```
ip prefix-list filter1 permit 129.0.0.0/8 le 24
```

Permit all routes:

```
ip prefix-list filter1 permit 0.0.0.0/0 le 32
```

Replace a prefix list entry

Remove entry `ip prefix-list filter1 permit 104.60.0.0/15` and replace with `ip prefix-list filter1 permit 198.0.0.0/8`:

```
no ip prefix-list filter1 permit 104.60.0.0/15 ip prefix-list  
filter1 permit 198.0.0.0/8
```

Manually pre-configured multi-AS VPNs

You can use Service Activator to manage manually pre-configured multi-AS VPNs.

When managing multi-AS VPNs with Service Activator, the domain-level property **Configure iBGP Peering** must be de-selected in the user interface. For more information, see the *Configuring VPN Services* guide.

PE routers in the same AS

iBGP peering must be manually pre-configured between PE routers that reside in the same AS.

For example:

```
router bgp 65057
  no synchronization
  no auto-summary
  neighbor 10.52.0.1 remote-as 65057
  neighbor 10.52.0.1 update-source 10.52.20.1

  address-family vpnv4 unicast
    neighbor 10.52.0.1 activate
    neighbor 10.52.0.1 next-hop-self
    neighbor 10.52.0.1 send-community extended
  exit
```

Inter-AS PE routers

eBGP peering must be manually pre-configured between PE routers where each PE router resides in a different AS.

For example:

```
router bgp 65057
  no synchronization
  no auto-summary
  neighbor 10.52.0.1 remote-as 65056
  neighbor 10.52.0.1 ebgp-multihop
  neighbor 10.52.0.1 update-source 10.52.20.1
```

```
address-family vpnv4 unicast
  neighbor 10.52.0.1 activate
  neighbor 10.52.0.1 next-hop-self
  neighbor 10.52.0.1 send-community extended
exit
```


Chapter 5

Configuration of MPLS VPNs

This chapter explains how Service Activator configures MPLS VPNs on Juniper E-series devices. It includes the following:

- A summary of the pre-requisites for configuring MPLS VPNs
- Configuring VRF tables and route targets
- Configuring PE-PE peering
- Configuring PE-CE routing using eBGP, RIP and static routing

Pre-requisites

The Juniper E-series device driver configures the PE routers that define the membership of a VPN. The information set up on each PE router defines the VPNs to which connected sites belong and the routes to and from these sites that are to be distributed throughout the VPN.

Service Activator does not configure the CE routers or the provider core routers.

Before setting up VPNs you should ensure the following:

- All routers are appropriately configured
- Domain-level parameters are appropriately set
- All routers and their interfaces within the VPN are correctly assigned roles

Pre-configuration of routers

Some pre-configuration of PE and P routers is required. For example, MPLS must be enabled. For full details of the preconfiguration required, see [Manual Pre-configuration on page 21](#).

Domain-level parameters

A number of BGP parameters may be set up at domain level on the **VPN BGP, ASN** and **VPN MPLS** property pages of the **Domain** dialog box:

- You can specify whether you want Service Activator to set up iBGP peering on the PE devices. See [Co-existence with previously configured iBGP on page 42](#).

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.

If Route Reflectors are used, iBGP peering must be deselected.

If Service Activator is to manage multi-AS VPNs, iBGP and eBGP peering must be configured on devices and Service Activator's configure iBGP peering capability must remain deselected. See [Manually pre-configured multi-AS VPNs on page 28](#).

- Set up the ASN for the domain (set on the **ASN** property page of the **Domain** dialog box).

If there is no ASN already configured on the device, Service Activator configures the device with the ASN specified in the user interface. If an ASN is already configured on the device, Service Activator ignores the ASN specified in the user interface and uses the one found in the configuration instead. This enables Service Activator to support multi-AS VPNs.

- You can enable **Allow AS in** which allows PE devices to re-advertise route prefixes containing one or more instances of the same ASN in the AS_PATH attribute. You specify the maximum number of instances allowed for an incoming prefix to be permitted by the PE device. The PE device denies incoming prefixes having more than the number of instances specified. See [Allow AS in on page 44](#).
- You can specify whether a **Global ASN** applies to sites within VPNs, or whether the ASN is **Set at site**. If a Global ASN applies, you can also enable **AS Override** which allows PE devices receiving route prefixes from the core, whose AS_PATH attributes have ASNs matching the ASN of their neighboring CEs, to substitute those ASN instances with the ASN of the service provider network. Prefixes with the substituted ASNs are then re-advertised to neighboring CEs. This is enabled by default. For more information, see [AS override on page 44](#).
- You can enable load-balancing between eBGP peers by setting a value for **Maximum Paths**. This controls the number of alternative routes to a given prefix that are maintained in a device's routing table. By default, this option is disabled and no alternatives are held. To enable load-balancing, you specify the number of routes that are maintained. For more information, see [eBGP load sharing on page 46](#).
- You can define the community attributes (Extended and/or Standard) that are used for routes distributed from PE devices, both for PE-PE peering and PE-CE. See [PE-PE community attributes on page 42](#) and [PE-CE community attributes on page 45](#).
- You can specify that the identity of iBGP peers and the integrity of data exchanged during iBGP sessions is to be verified using MD5 Authentication. See [MD5 authentication on page 42](#).
- You can specify which interface is configured as the loopback interface on devices in this domain. The default is 0.
- Service Activator supports multiple Loopback0 addresses. The device driver has been altered to use ERX style configuration for interfaces. An add-on component allows the creation of loopback0 addresses for each VRF.

For full information, see the *Configuring VPN Services* guide.

Discovery and role assignment

In an MPLS domain, the core provider network is assumed to use public addresses. All CE routers are assumed to use private addresses. An IP address or DNS name must be specified in order to discover all devices in the domain.

All devices within the network must be assigned the correct system-defined roles (that is, PE routers must be classified as Gateway devices, P routers as Core devices and CE routers as Access devices). Interfaces to be configured must also be assigned the correct roles. You can assign user-defined roles as well as the system-defined roles. The recommended way of assigning roles is by means of role assignment rules, which automatically assign roles during device discovery.

For full information on role assignment rules, see the *Network Discovery and Basic Setup* guide.

Configuring VRF tables and route targets

The Juniper E-series device driver configures the appropriate VRF (VPN Routing/Forwarding Instance) tables and associated route targets on the PE devices.

Each customer site connects to a PE interface or sub-interface. This interface is assigned to a VRF table, which defines the VPN membership of a customer site.

VRF tables hold routing information that defines how packets from a given site are routed across one or more VPNs to other sites. They are private routing tables containing IPv4 routes that have been learnt from CE routers. They do not form part of the PE router's own routing tables.

On the Service Activator user interface, data specific to the customer site and the PE interface are set up on the Site dialog box. Data relating to the VPN and its connectivity are set on the VPN dialog box.

VRF tables

The Juniper E-series device driver configures the appropriate VRF (VPN Routing/Forwarding Instance) tables and associated route targets on the PE devices.

E-series devices use virtual routers to implement VRF tables. The following command allows you to access a specified virtual router. Subsequent commands apply to the specified virtual router.

```
virtual-router virtual-router-name
```

The following command is configured at the root level. The *vrf-name* is generated automatically unless a user-specified name has been defined.

```
ip vrf vrf-name
```

The following command is implemented on the appropriate interface to associate it with the specified VRF table:

```
ip vrf forwarding vrf-name
```

By default, Service Activator generates VRF table names of the form *Orch_RD-number*. However, you can define specific names for VRF tables on selected interfaces if you do not want to use the system-calculated ones.

If you specify a VRF table name you must ensure that the name you enter is not the same as the name of any other virtual router configured on the device.

Note that on older Juniper E-series devices, VRF table names are limited to 15 alphanumeric characters. Starting with Service Activator version 4.0, the new 32 character VRF table name length limit is supported.

If you are using Juniper E-series devices which support only the 15 character limit with Service Activator 4.0 or greater, any name that is longer than 15 characters will fail to instantiate on the router. To remedy this, shorten the name to 15 characters or less.

Do not modify VRF parameters of an unmanaged device. Service Activator may get out of sync with a device if VRF parameters are manually changed while a device is unmanaged, and the device is then re-managed. In this event, the device may not have correct VRF configuration.

If this occurs, manually remove the incorrect VRF configuration from the device.

When a Service Activator supplied route distinguisher on a VRF conflicts with a manually supplied route distinguisher on a VRF, once you modify the manually supplied route distinguisher Service Activator may still report a conflict even though none exists.

If a user-defined VRF table name begins with a digit, Service Activator prepends a `'_'` character to the name.

By default, Service Activator automatically generates a site-specific VRF table name for each site that participates in a VPN. However, if you wish to apply the same RD number to all sites that participate in the VPN, the same VRF name will also apply (auto-generated or user-defined). See [RD number per VPN on page 36](#).

Route distinguishers

Customer networks typically use private addresses. Addressing overlaps between customers may occur when they connect to the public Internet or to the provider's NOC. To avoid this problem, iBGP prefixes a site identifier, known as a route distinguisher or RD number, to each route associated with a particular site. This ensures that VPN routes are unique within the Internet.

The new route is part of the VPN-IPv4 address family – a BGP address family added as an extension to the BGP protocol.

The RD number can be in either of the following formats:

32-bit IP address: 16-bit number

16-bit ASN number:32-bit number

Service Activator normally generates RD numbers automatically, using the ASN for the high-order-no and the unique system ID of the Site object for the low-order-no. For example:

1:3125

However you can override these defaults and specify your own RD numbers if you wish. The following command is configured at the VRF level:

```
rd rd-number
```

Note that Service Activator does not validate that manually-generated RD numbers are unique, although system-generated RDs are always unique. This is to ensure compatibility with user-defined configuration.

RD number per VPN

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

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.

On the user interface, these settings are specified on the **VRF** page of the **VPN** dialog box.

In addition, you can override the name generation rules so that if a site is part of only one VPN, the VRF table name and RD are derived from the VPN. If a site is part of multiple VPNs, the VRF table name and RD are derived from the site.

If a single RD number/VRF table name is set per VPN, the settings for VRF re-use/reduction must also be set at VPN level. See [VRF re-use/reduction on page 39](#).

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.

VPN topology and route targets

The connectivity of the VPN can be one of the following:

- **Mesh** – all sites have connectivity to all other sites
- **Hub and Spoke** – one or more hub sites has access to all other sites; spoke sites can access the hub only
- **Management** – works in the same way as hub and spoke, but is used to ensure connectivity to CE devices.

When setting up a VPN, you have to set its connectivity, and for a hub and spoke or management VPN, select the hub site(s).

To create a fully-meshed VPN, each site's VRF table imports and exports the same routes. In a hub and spoke or management VPN the VRF table at the hub site imports routes from all other hub sites and all spoke sites, and exports routes to other hub sites and to the spoke sites. VRF tables at spoke sites export routes only to the hub site and import routes only from the hub site.

A route target (RT) identifies a set of sites within a VPN to which a PE device distributes routes.

Route targets are used to create the VPN topology. Each VPN must have a unique route target number.

The RT is implemented as a BGP extended community. A BGP community groups a set of destinations that share a common property – in this case, a set of routes that are to be distributed to a set of CE sites. The RT is added to the route by the ingress PE device and used by the egress PE device to determine whether a received route is destined for a VPN that the PE services.

Service Activator creates one or more BGP communities per VPN, depending on the VPN topology:

- If the VPN is fully-meshed, Service Activator creates one community – every site receives routing information from all other sites
- If the VPN is a hub and spoke or management VPN, Service Activator creates two communities

In a hub and spoke topology, there are effectively two ‘sets’ of devices – one set that consists of the hub site or sites and another set consisting of the spoke sites. Routes from the spoke sites are only distributed to the hub site(s), routes from the hub site(s) are distributed to all spoke sites and imported by all other hub sites.

The RT number can be in either of the following formats:

32-bit IP address: 16-bit number

16-bit ASN number:32-bit number

Service Activator normally generates RT numbers automatically, using the ASN for the high order number and the unique system ID of the VPN for the low order number.

For example:

20:4926

In a hub and spoke VPN topology, Service Activator generates two RT numbers – one for the hub site(s), generated as indicated above, and one for all spoke sites, generated by incrementing the hub low order number by 1.

If you wish, you can specify your own RT numbers for hub, spoke or fully-meshed sites within a VPN if you do not want to use the system-generated default values. You can easily reassign RT numbers to sites within a VPN, if for example, it has been imported from a different system or it is to be exported to a different system.

You can specify any number of RT values per VPN and specify whether a value is used for VRF import, VRF export, or neither for hub, spoke and fully-meshed behaviors.

The `route-target` command configures a PE’s VRF table with an import and an export policy which allow RT values to be specified for VRF import and VRF export. Routes whose iBGP VPNv4 extensions have RT values matching those in the import policy are imported to the VRF table. All routes exported from the VRF table will have their VPNv4 extensions attached with RT values specified in the export policy ensuring that these private routes are only advertised to neighboring PE routers that share the same VPNs.

The following commands are configured at the VRF level:


```
route-target export rt-number
```

This command defines the routing information to be exported to the target VPN. The RT number identifies the VPN to which routes are exported.

```
route-target import rt-number
```

This command defines the routing information to be imported from the target VPN. The RT number identifies the VPN from which routes are imported.

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.

VRF re-use/reduction

A VRF table is set up on the device for each PE interface that is a member of a VPN. However, if multiple VRF tables contain exactly the same routes (for example if one site connects to two interfaces, or there are two sites that are members of the same VPN) Service Activator will normally reduce them to just one, in order to minimize resource usage. This is known as VRF re-use or VRF reduction.

In some cases automatic VRF re-use may not be required. For example, you may want to provision dual links to customer sites in order to implement load balancing, requiring a separate VRF table for each connecting interface, or to reduce the impact of future re-configuration. In this case you can override VRF re-use by specifying that particular interfaces are always to have their own VRF table, and by specifying that other VRF tables are allowed to be merged with this VRF table by selecting the Shareable option.

On the user interface, the **Force Install** and **Shareable** options can be selected per interface on the **VRF** property page of the **Site** dialog box.

Note that if you are setting up a single VRF table name/RD number per VPN, settings for VRF re-use are made at the VPN level. In this case it is not possible to select Force Install. See [RD number per VPN on page 36](#).

Co-existence with pre-defined VRF tables

If an MPLS VPN has already been manually configured on a network, Service Activator is able to 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 several parameters which can be defined per site or per VPN.

Site or VPN property	Manually pre-configured VRF		
	No control	Control and preserve content	Control and 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 are preserved.

Service Activator’s normal behavior is to maintain pre-defined VPN configuration. However, to automatically remove all VPN configuration on PE routers that are not controlled by Service Activator, you can use the `-DisableVpnPreservation` command-line parameter.

Previously-defined export maps

You can apply a user-defined export map to the export policy configured by Service Activator. The export map exports only the VRF table routes whose prefixes match those specified in the export map to other PE devices. The export map tags these routes with only the RT numbers of sites that need to receive those routes.

The name of the pre-existing export map must be specified to Service Activator.

For details of setting up an export map, see [Pre-defined export maps on page 24](#).

Configuring PE-PE peering

In order to exchange information throughout the VPN, each PE router needs to run an iBGP session with each other PE router connected to a site within the same VPN.

By default, this configuration is set up automatically.

Configuring iBGP

iBGP is the protocol used for communication of VPN routes between PE devices in an MPLS VPN. In order for devices to exchange routing information, an iBGP session must be configured between the PE devices that comprise the VPN.

If the domain-level parameter **Configure iBGP Peering** is selected on the **VPN BGP** property page of the **Domain** dialog box, Service Activator configures adjacencies between PE devices depending on the VPN's topology:

- For fully-meshed VPNs, Service Activator creates a full mesh of iBGP adjacencies
- For hub and spoke or management VPNs, Service Activator configures iBGP peering between each spoke and the hub site(s)

The following command is configured at the VRF level to configure iBGP:

```
router bgp asn
```

The following commands are entered at the bgp level:

```
neighbor ip-addr remote-as asn
```

This command is required for each PE router with which the router will peer. It specifies that the PE router is a member of the BGP routed network and identifies the address to which routing updates should be sent. The IP address is the loopback address of the PE router, and the ASN is the ASN of the core network.

```
neighbor ip-addr update-source loopback-if
```

This command tells BGP that the PE's loopback interface is used for the iBGP neighbor TCP connection. The IP address is the IP address of the neighbor PE.

```
neighbor ip-addr next-hop-self
```

Advertises the local PE router as the next hop for an iBGP peer.

The following command is configured at the root.bgp.vpnv4 level:

```
neighbor ip-addr activate
```

Activates iBGP peering in the vpnV4 address-family.

PE-PE community attributes

You can specify that routes advertised to the neighbor CE router contain the standard community attribute as well as the extended community attribute which is configured by default. The following commands are configured at the bgp level:

```
neighbor ip-addr send-community standard
```

```
neighbor ip-addr send-community extended
```

The extended community attribute, which includes Route Target and Site Of Origin, is always configured and should not be removed. The standard community attribute is optional. If both community attributes are required, both commands are configured separately rather than the `send-community both` command.

Co-existence with previously configured iBGP

If iBGP peering is already installed on PE routers, for example if Route Reflectors are used, the existing configuration can be preserved. On the Service Activator user interface, this is controlled by deselecting the **Configure iBGP Peering** option at the domain level.

Note that if **Configure iBGP Peering** is deselected, the system will leave all iBGP configuration on the device untouched. Service Activator will not configure any iBGP commands, so you should ensure all configuration is correct.

MD5 authentication

The identity of iBGP peers and the integrity of data exchanged during iBGP sessions can be verified using MD5 Authentication. This option uses the MD5 digital signature algorithm and a specified key of up to 255 characters to generate a checksum of the iBGP data that is to be sent from a PE device to its peer. The iBGP data and its checksum are then sent to the peer device using TCP. The recipient device uses MD5 and the same key to generate a checksum of the received iBGP data. If both checksums match, the identity of the sender and the integrity of the received iBGP data is verified.

If this option is selected, the following command is configured at the bgp level:

```
neighbor ip-addr password private-key
```

PE-CE configuration using eBGP

In order to exchange information to and from customer sites in the VPN, each PE router also needs to communicate with each of its external neighbors – the CE

routers to which it is connected. Service Activator supports the eBGP routing protocol. Static routes can also be defined (either alone or in conjunction with eBGP and/or RIP).

The effect is to advertise network reachability information between the CE and the PE, which in turn converts IPv4 addresses to VPN-IPv4 addresses for traffic passing from the CE to the PE and vice versa.

The details here explain the configuration of the PE routers. The corresponding configuration of the CE routers is not performed by Service Activator. See [CE routers on page 23](#).

If eBGP is used, you need to specify:

- The ASN of the site
- The IP address of the corresponding interface on the CE router
- The number of times the same ASN can appear in the AS_PATH of an incoming prefix for it to be accepted by the site PE or all PEs in the domain

You can optionally specify:

- The number of times the same ASN can appear in an incoming prefix for it to be accepted by the site PE or all PEs in the domain
- AS Override for the site
- Authentication for a PE-CE session
- Send community preferences for the site
- Prefix limit for the site
- Prefix filter for the site
- Where multiple PE interfaces are associated with a site, you can set the local preference for each interface.

On the user interface, eBGP configuration is controlled by settings on the **EBGP**, **EBGP Adv** and **EBGP Damp** pages on the **Site** dialog box.

Service Activator configures the following E-series commands automatically:

```
address-family ipv4 vrf vrf-name
```

Enters the IPv4 address family submode for configuring routing protocols. Subsequent commands affect eBGP routing in the context of the named VRF table only. The VRF table name matches that in the `ip vrf forwarding` command run on the specific interface.

```
neighbor ip-addr remote-as asn
```

Defines an eBGP session to exchange information with the neighboring CE router. The IP address identifies the access interface on the CE router and the ASN identifies the autonomous system of which the CE router is a member.

```
neighbor ip-addr description site-name
```

```
neighbor ip-addr update-source PE-CE-interface
```

```
neighbor ip-addr activate
```

These commands are required for each PE router. They enable the exchange of information with neighboring PE routers. The `update-source` command tells BGP that the PE's loopback interface is used for the eBGP neighbor TCP connection. The `description` command is added to associate the set of commands with a site name.

All VPN-IPv4 addresses are then advertised to the neighbor router.

Allow AS in

You can specify the maximum number of times the same ASN is allowed to occur in the `AS_PATH` attribute of a route prefix advertised to the PE device for the prefix to be permitted and then re-advertised to neighboring CEs by the PE device. The value can be from 0 to 10; the default is 0.

Within Service Activator, the Allow AS in value can be specified at domain level, to apply to all sites, or set up for individual sites.

If this is selected, the following command is configured:

```
neighbor ip-addr allowas-in n
```

AS override

You can specify that the ASN of a provider is used to override the ASN of a site. When AS override is turned on, a PE device that receives route prefixes whose `AS_PATH` attributes have one or more ASNs matching the ASN of its neighboring CE devices, substitute those ASN instances with the ASN of the service provider network. Prefixes with the substituted ASNs are then re-advertised to neighboring CE devices. The PE device also adds its ASN to routes before exporting them to the CE device.

This allows CE devices to accept routes that have been re-advertised by devices having the same ASN, and which would otherwise be rejected. Normally, a CE device rejects routes whose `AS_PATH` attribute contains ASNs matching its own ASN, to prevent routing loops.

Within Service Activator, AS override can be specified at domain level, to apply to all sites, or set up for individual sites.

If this option is selected, the following command is configured at the root.bgp.ipv4-vrf level:

```
neighbor ip-addr as-override
```

PE-CE community attributes

You can specify that routes advertised to the neighbor CE router contain the standard or extended community attribute or both.

The following commands are configured at the root.bgp.ipv4-vrf level:

```
neighbor ip-addr send-community standard
```

```
neighbor ip-addr send-community extended
```

Note that if both community attributes are selected, both commands are configured separately rather than the `send-community both` command.

Authentication

The identity of eBGP peers and the integrity of data exchanged during eBGP sessions can be verified using Authentication. If this option is selected, the following command is configured at the ipv4-vrf level:

```
neighbor ip-addr password private-key
```

Local preference

Where multiple PE interfaces are associated with a site, the local preference for an interface can be set. The preference value may be between 1-4294967295, and the higher the value the higher the priority. The default is Router Default (100). Local preference is configured by means of a route map, which can include other conditions.

In the Service Activator user interface, the local preferences parameter is set on the **EBGP** page of the **Site** dialog box.

If this option is selected, the following command is configured at the IPv4 level:

```
neighbor ip-addr route-map route-map-name in
  route-map route-map-name permit sequence-number
  set local-preference value
```

Site of origin

Site of Origin (SOO) is configured automatically for sites that have more than one CE to PE connection. It identifies the site from which the PE router learned the route and prevents routing loops from occurring when a site is multi-homed. SOO is configured by means of a route-map.

```
neighbor ip-addr route-map route-map-name in
  route-map route-map-name permit sequence-number
  set community extended soo soo-id
```

The unique ID is automatically generated by Service Activator.

Route prefix limits and filters

If a prefix list file is specified, routes whose prefixes match those in the prefix list will either be allowed or rejected by the PE router depending on their designation in the prefix list.

You can specify a maximum number of eBGP IP route prefixes that can be received by the PE router from its CE neighbor. You can specify a threshold percentage of the maximum at which the PE router generates a warning log message. The PE router terminates the peering if the maximum is exceeded. Alternatively you can request that a warning is issued, but routes are still accepted and the peering is not terminated.

In the Service Activator user interface, the prefix limits and filter parameters are set on the **eBGP Adv** page of the **Site** dialog box.

If this option is selected, the following commands are configured

```
neighbor ip-addr prefix-list list-name in
neighbor ip-addr prefix-list list-name out
```

The in and out arguments specify that the prefix list file applies to either incoming routes (CE-PE) or outgoing routes (PE-CE).

```
neighbor ip-addr maximum-prefix maximum threshold
neighbor ip-addr maximum-prefix maximum warning-only
```

eBGP load sharing

You can enable load-balancing between eBGP peers. This allows BGP to select more than one eBGP path to a given prefix over which traffic can be shared. Routes for each path are maintained in a PE device's routing table. By default, this option is disabled and no alternative routes are held.

In the Service Activator user interface, the eBGP maximum paths parameter is set at global level on the **VPN BGP** page of the **Domain** dialog box and overridden for specific devices by a setting on the **EBGP Adv** page of the **Site** dialog box.

If this option is selected, the following command is configured at the IPv4 level:

```
maximum-paths eibgp max-paths
```

where:

max-paths specifies the maximum number of paths allowed, in the range 2-6

Note that multi-path load sharing can affect VRF reduction – there is no VRF reduction if different Maximum Paths values are set on different interfaces.

Route dampening

Route dampening is a mechanism that attempts to minimize network instability by suppressing the advertisement of poorly-behaved routes. Penalties are applied when a route is withdrawn, readvertised 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 readvertised.

If this option is selected, the following command is configured at the bgp level:

```
bgp dampening half-life reuse suppress max_suppress_time
```

where:

half-life Time, in minutes, at which a penalty applying to a route is decreased by half. Range is 1-45.

reuse When the penalty applying to a route falls below this value, the route is unsuppressed. Range is 1-20000.

suppress A route is suppressed when its penalty exceeds this limit. Range is 1-20000.

max_suppress_time The maximum time, in minutes, that a route can be suppressed. Range is 1-20000.

Route redistribution into eBGP

Service Activator allows you to control the redistribution of routes into eBGP from static routing and directly-connected networks.

Redistributing routes between protocols brings with it the risk of introducing routing loops and convergence problems. However, you can filter and refine the redistribution of routes by associating a manually pre-configured route map with redistributed routes.

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

Where a site is multi-homed, you can specify route distribution metrics per interface.

In the user interface, route redistribution is set up on the **Redist** page of the **Site** dialog box. You can specify the protocol-specific metric to apply to static routes redistributed into eBGP, and define a manually pre-configured route map to apply to static routes redistributed into eBGP.

Depending on the route redistribution options selected, the following commands can be configured by Service Activator:

```
redistribute static metric metric route-map route-map-name
```

Redistributes static routes into BGP.

```
redistribute connected metric metric route-map route-map-name
```

Redistributes connected routes into BGP.

```
default-information originate
```

Specifies that the default route (0.0.0.0) is distributed via iBGP to peers within the VPN. Without this command, BGP will not propagate the default route.

PE-CE configuration using RIP

This section describes the commands configured if RIP is used as the PE-CE routing protocol. The corresponding configuration of the **CE** routers is not performed by Service Activator. See [CE routers on page 23](#).

As well as RIP, Service Activator supports eBGP and OSPF routing protocols and static routing. See [PE-CE configuration using eBGP on page 42](#) and [PE-CE configuration using static routing on page 51](#).

Note: Once RIP has been configured on a Juniper E-series device, it remains configured even if RIP is removed from all interfaces on the device and all other RIP parameters are removed. In other words, the command `router rip` will remain on the device. If required, remove it manually.

When configuring RIP between PE and CE routers in an MPLS VPN, Service Activator takes different approaches for routers which support the Default Passive Interface feature, depending on the existing RIP configuration (if any) on the device.

When no RIP configuration exists on the router, or it is set up for all interfaces to be passive by default:

- In the case where no RIP configuration exists on the router, Service Activator sets up all interfaces to be passive on the router using the `passive-interface default` command. This puts all interfaces into passive mode unless specifically configured otherwise.
- If the `passive-interface default` command already exists on the router, this configuration is left untouched. Again, all interfaces are put into passive mode unless specifically configured otherwise.
- For specific interfaces in Service Activator, if **Passive Interface** is unchecked on the **Site** dialog box, **RIP** property page, the interface is configured with the `no passive-interface` command to turn off passive status. If **Passive Interface** is checked, no change is made — the default passive status for the interface is maintained.

When the router is set up for all interfaces to be active by default:

- In the case where the RIP is configured on the router and interfaces are active by default, this configuration is left untouched. This puts all interfaces into active mode unless specifically configured otherwise.
- For specific interfaces in Service Activator, if **Passive Interface** is unchecked on the **Site** dialog box **RIP** property page, no change is made — the default active status for the interface is maintained. If **Passive Interface** is checked, the interface is configured with the `passive-interface` command to enable passive status for that interface.

When no RIP configuration exists on the router, or it is set up for all interfaces to be passive by default, Service Activator installs the following configuration:

```
router rip
    Configures RIP on the specified router.
passive-interface default
```

Makes all interfaces passive by default (i.e. disables the sending of routing updates on all interfaces.)

```
network address
```

Specifies addresses of interfaces that run RIP.

```
no passive-interface {interface-type interface-number}
```

For interfaces on sites which have **Passive Interface** is unchecked on the **Site** dialog box **RIP** property page, over-rides the default setting to make the specific interface non-passive (i.e. enables the sending of routing updates on the specified interface.)

When the router is set up for all interfaces to be active by default, Service Activator installs the following configuration:

```
router rip  
version 2
```

Configures RIPv2 on the specified router.

```
network address
```

Specifies addresses of interfaces that run RIP.

```
passive-interface {interface-type interface-number}
```

For interfaces on sites which have **Passive Interface** is checked on the **Site** dialog box **RIP** property page, over-rides the default setting and specifies that the interface will not run RIP, even if selected by a network statement (i.e. disables the sending of routing updates on the specified interface.).

Route redistribution into RIP

Service Activator allows you to control the redistribution of routes into RIP from eBGP, static routing and directly-connected networks.

Note: Redistribution into OSPF is not valid on Juniper E-series devices.

Redistributing routes between protocols brings with it the risk of introducing routing loops and convergence problems. However, you can filter and refine the redistribution of routes by associating a manually pre-configured route map with redistributed routes.

Where a site is multi-homed, you can specify route distribution metrics per interface.

Route redistribution is set up on the **Redist** property page of the **Site** dialog box. You can specify the protocol-specific metric and define a manually pre-configured route map to apply to routes redistributed into RIP.

Depending on the route redistribution options selected, the following commands can be configured by Service Activator:

```
redistribute bgp asn metric n route-map route-map-name
redistribute static metric n route-map route-map-name
redistribute connected metric n route-map route-map-name
```

Redistributes BGP, static routes and connected routes respectively into RIP. The metric value for BGP routes defaults to 1, and for connected and static routes the metric defaults to 0.

PE-CE configuration using static routing

Static routes can be configured in conjunction with eBGP or RIP routing or used alone to define routing between the PE and CE. If static routing is required, you need to define the destinations that can be reached from each PE interface, that is, the CE device (the address of the loopback interface) and any connected networks.

In the Service Activator user interface, static routes are engaged on the **Connectivity** property page of the **Site** dialog box and set up on the **Static Routing** property page.

On the **Static Routing** property page of the **Site** dialog box, the **Global** checkbox does not function on Juniper E-series devices (It normally specifies that the next-hop-address is an address that is in the routing table and not in the VRF table). The **Permanent** checkbox (which normally specifies that the static route will not be removed, even if the interface shuts down) does not function on Juniper E-series devices.

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. This controls the following command:

```
redistribute static metric n route-map route-map-name
```

Configuring how the Next Hop parameter is specified

A number of choices are available on the **Static Routing** property page of the **Site** dialog box to control the way the next hop value will be specified in the `ip route vrf` command in the static route. From the **Next Hop** dropdown, choices are:

IP Address & I/F: specify next hop value using both the IP Address and the interface name

IP Address Only: specify next hop value using only the IP Address. This is applicable, for example, to configurations such as MLPPP where load sharing is needed over a set of interfaces.

I/F Only: specify next hop value using only the interface name. This can be used when configuring point-to-point links.

Null0I/F: specify the Null0 interface as the next hop. This filters out all traffic going to the specified destination address and drops it.

Configuring a Null0 static route

You can specify the Null0 interface as the Next Hop of a static route by selecting **Null0I/F** for the Next Hop parameter. See [Configuring how the Next Hop parameter is specified on page 52](#).

Static Routing Commands

The relevant commands for configuring static routing on the PE router are configured automatically by Service Activator.

An `ip route` command is included for each static route defined. It has the following format:

```
ip route vrf vrf-name ip-addr mask next-hop interface [distance] [tag tag]
```

where:

<i>vrf-name</i>	Matches that in the <code>ip vrf forwarding</code> command run on the specific interface.
<i>ip-addr</i>	The IP address of the defined static route.
<i>mask</i>	The mask of the defined static route.
<i>next-hop</i>	IP address of the next hop that can be used to reach the destination.
<i>interface</i>	Interface on which the VRF table is configured.

<i>distance</i>	Cost metric of the route (1-255). Specifies the weighting of the route entry in the VRF IP routing table.
<i>tag</i>	Specifies a value in the range 1 to 4294967295 to identify the static route, allowing it to be used by route-map match statements controlling the redistribution of routes.

It is possible for Service Activator to co-exist with manually-configured static routes which apply to VRFs not under the control of Service Activator. However, in some circumstances, such as VRF reduction, the system can remove manually-configured static routes or non-supported parameters.

Commands reflecting the various Next Hop dropdown choices are as follows.

IP Address & I/F:

```
ip route vrf vrf-name prefix mask [next-hop-address] [interface {interface-number}] [distance] [tag tag]
```

IP Address Only: All fields and options are valid

```
ip route vrf vrf-name prefix mask [next-hop-address] [distance] [tag tag]
```

I/F Only:

```
.ip route vrf vrf-name prefix mask [interface {interface-number}] [distance] [tag tag]
```

Null0:

```
ip route vrf vrf-name prefix mask null0 [distance] [tag tag]
```


Chapter 6

Troubleshooting

This chapter provides hints and tips for diagnosing and fixing problems occurring with the Juniper E-series device driver. It includes the following:

- Monitoring the device driver logs
- Communication problems and how to fix them
- Useful commands for checking configuration

Checking the Service Activator device audit logs

Each installed instance of the Juniper E-series device driver records all device configuration changes that it makes in a specific device audit log file. You can check these log files to see if configuration is being successfully applied to a device.

A new log file is created each day, the first transaction after midnight resulting in a new log file being created. The same log file is used if a device driver is stopped and started within a 24-hour period.

- On Solaris systems, the log files are in the **AuditTrails** directory of the host running the device driver. By default this is **/opt/OracleCommunications/ServiceActivator/AuditTrails**
- On Windows systems, the log files are created in the **AuditTrails** directory of the host running the device driver. By default this is **Program Files\ Oracle Communications\IP Service Activator\AuditTrails**

Each file is named **<day>.unisphere.audit.log**, e.g. **Tue.unisphere.audit.log**. After a week, the log files are automatically overwritten, so you should archive them within this period if you want to keep them.

Log files are text files, recording the date, time and details of each configuration change made to the devices controlled by the device driver.

The following show example commands contained in an audit log:

```
Fri 19/10/01 17:07:42 UTC|10.0.0.42:PE2|*** Begin Transaction ***
Fri 19/10/01 17:05:23 UTC|10.0.0.42:PE2|Start Configuring
Fri 19/10/01 17:05:24 UTC|10.0.0.42:PE2|router bgp 1
Fri 19/10/01 17:05:24 UTC|10.0.0.42:PE2|address-family vpnv4
Fri 19/10/01 17:05:25 UTC|10.0.0.42:PE2|no neighbor 10.1.2.0 send-
community
Fri 19/10/01 17:05:25 UTC|10.0.0.42:PE2|no neighbor 10.1.2.0 activate
Fri 19/10/01 17:05:25 UTC|10.0.0.42:PE2|no neighbor 10.1.2.0 next-hop-
self
Fri 19/10/01 17:05:25 UTC|10.0.0.42:PE2|exit
Fri 19/10/01 17:05:25 UTC|10.0.0.42:PE2|no neighbor 10.1.2.0 update-
source 10.1.2.2
Fri 19/10/01 17:05:26 UTC|10.0.0.42:PE2|no neighbor 10.1.2.0 remote-as
1
Fri 19/10/01 17:05:26 UTC|10.0.0.42:PE2|neighbor 10.1.2.1 remote-as 1
Fri 19/10/01 17:05:26 UTC|10.0.0.42:PE2|address-family vpnv4
```

```
Fri 19/10/01 17:05:26 UTC|10.0.0.42:PE2|neighbor 10.1.2.1 activate
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|neighbor 10.1.2.1 send-
community
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|exit
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|neighbor 10.1.2.1 update-source
10.1.2.2
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|address-family vpv4
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|neighbor 10.1.2.1 next-hop-self
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|exit
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|exit
Fri 19/10/01 17:05:27 UTC|10.0.0.42:PE2|ip vrf Orch_1_4262
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|no route-target import 1:2550
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|no route-target export 1:2550
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|route-target import 1:5231
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|route-target export 1:5230
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|exit
Fri 19/10/01 17:05:28 UTC|10.0.0.42:PE2|Exiting
Fri 19/10/01 17:05:29 UTC|10.0.0.42:PE2|*** Commit Transaction ***
```

Communication problems

If you are unable to communicate with a router, check the following:

- Ensure that enable passwords are correctly set up, using the `enable password` or `enable secret` commands.
- Ensure a password for the router's VTY line is set up (check for a password command following the `line vty 0 4`)
- Ensure that an SNMP Read community is set up on the device, e.g. the following command is included:

```
snmp-server community public RO
```

If a community other than "public" is used, you need to amend the appropriate parameter within Service Activator.

Useful E-series commands

There are a number of Juniper E-series commands that can be useful in checking device configuration. For full details of the command syntax and an explanation of the reported information, see the Juniper E-series documentation.

General configuration

The following commands report general status and configuration:

```
show processes cpu
```

Displays detailed CPU statistics, allowing you to monitor router performance.

```
show config
```

Displays your system's entire configuration, including information for all virtual routers.

The results will show global configuration first, and then configuration for each virtual router in turn.

```
show config virtual-router virtual-router-name
```

Lists the current router configuration for the *virtual-router-name* entered. The command will not show any child virtual routers or global configuration.

```
show version
```

Indicates the current operating system version and general information about the router.

Some commands issued on the main router which cause interaction between a main virtual router and one of its VRF children don't appear in the configuration of the main router. Examples include setting up static routing and `ip vrf forwarding vrf-table-name` on interfaces. Commands appear on the virtual router that represents the VRF table, sometimes in a slightly different format.

VPN configuration

The following commands are useful for checking the configuration of MPLS VPNs.

Before entering a command, first go to the virtual router you require by entering the command `virtual-router virtual-router-name`.

show ip bgp neighbor

This command shows information about the BGP neighbor of the device. On the PE and CE routers, look for lines starting:

```
BGP state =
```

Alternatively use modifiers such as `:include BGP` to show only those lines including BGP.

There will be one line for each peer device. A state of Established means the VPN is functioning correctly. If it is not established, this could indicate that the ASN or neighbor IP address has not been set correctly on the property pages of the **Site** dialog box. Alternatively, tag switching may not be running. An example is as follows:

```
UNI:PE1#s ip bgp nei
```

```
BGP neighbor is 172.16.1.3, remote AS 1, internal link
```

```
BGP version 4, remote router ID 172.16.1.3
```

```
BGP state = Established, up for 00:56:03
```

```
Last read 00:00:04, hold timeUNI:PE1# is 180, keepalive interval is 60 seconds
```

```
Neighbor capabilities:
```

```
Route refresh: advertised and received
```

```
Address family VPNv4 Unicast: advertised and received
```

```
Received 60 messages, 0 notifications, 0 in queue
```

```
Sent 60 messages, 0 notifications, 0 in queue
```

```
Route refresh request: received 0, sent 0
```

```
Minimum time between advertisement runs is 5 seconds
```

show ip route

This command shows the IP routes on the device. An example is as follows:

```
UNI:PE1#show ip route
```

```
Protocol/Route type codes:
```

```
I1- ISIS level 1, I2- ISIS level2,
```

```
I- route type intra, IA- route type inter, E- route type external,
```

```
i- metric type internal, e- metric type external,
```

```
O- OSPF, E1- external type 1, E2- external type2,
```

N1- NSSA external type1, N2- NSSA external type2

Prefix/Length	Type	Next Hop	Dist/Met	Intf
10.1.2.0/32	Connect	10.1.2.0	0/1	loopback0
10.1.2.1/32	0-I	172.16.0.1	110/3	atm3/1.1
10.1.2.2/32	0-I	172.16.0.1	110/3	atm3/1.1
10.1.2.3/32	0-I	172.16.0.1	110/3	atm3/1.1
10.1.3.1/32	0-I	172.16.0.1	110/2	atm3/1.1
172.16.0.0/30	Connect	172.16.0.2	0/1	atm3/1.1
172.16.1.0/30	0-I	172.16.0.1	110/2	atm3/1.1
172.16.2.0/30	0-I	172.16.0.1	110/2	atm3/1.1
172.16.3.0/30	0-I	172.16.0.1	110/2	atm3/1.1

show ip route vrf vrf-name

Shows the routing tables specific to the named VRF, i.e. the VPNv4 table rather than the global routing table.

show mpls interface

This command indicates the interfaces on a device on which mpls is enabled.

show ip bgp

Displays a list of all BGP neighbors for the device.

show ip bgp vpnv4

Displays VPN address information from the BGP routing table.

show ip vrf

Shows the set of VRFs and associated route targets. For example:

```
D4500-1#show ip vrf
Name          Default RDInterfaces
FRED          10:1
Orch_598     1:1012   Serial1
```

show mpls forwarding

This command shows whether MPLS forwarding is operational on a particular device.

clear ip bgp

This command can be run on a PE device to reset BGP routing. If the wildcard is specified (`clear ip bgp *`), all BGP routing will be reset; alternatively a particular device can be specified. The device should start to relearn routes immediately, so you can use this command to clear BGP routing and then run the `show ip bgp` command to check if the routes are present.

clear ip route

This command can be run on a CE device to delete routes from the IP routing table. If the wildcard is specified (`clear ip route *`) all IP routing will be reset; alternatively a particular device can be specified. If BGP is running correctly, this will leave just BGP routes. You can run the `show ip route` command to check the routes.

Appendix A

Sample Device Configuration

This appendix provides example configuration of routers involved in an MPLS VPN. The following examples are given:

- Sample configuration for a spoke site
- Sample configuration for a three-spoke hub site

Sample configuration for a spoke site

```
virtual-router PE1

ip vrf Orch_1_4776
  rd 1:4776
  route-target export 1:5230
  route-target import 1:5231
!
interface loopback 0
  ip address 10.1.2.0 255.255.255.255
!
router bgp 1
  no synchronization
  neighbor 10.1.2.1 remote-as 1
  neighbor 10.1.2.1 update-source 10.1.2.0
!
  address-family vpnv4 unicast
  neighbor 10.1.2.1 activate
  neighbor 10.1.2.1 next-hop-self
  neighbor 10.1.2.1 send-community
  neighbor 10.1.2.1 send-community extended
  exit-address-family
!
  address-family ipv4 vrf Orch_1_4776
  no synchronization
  no auto-summary
  neighbor 192.168.2.2 remote-as 200
  neighbor 192.168.2.2 update-source 192.168.2.1
  neighbor 192.168.2.2 description s2.2
  neighbor 192.168.2.2 send-community
  neighbor 192.168.2.2 as-override
  exit-address-family
!
```

```
-----  
virtual-router PE1:Orch_1_4776  
  
interface atm 3/0.2 point-to-point  
  atm pvc 2 1 2000 aal5snap 0 0 0  
  ip address 192.168.2.1 255.255.255.252
```

Sample configuration for a 3-spoke hub site

```
virtual-router PE3  
  
ip vrf Orch_1_4268  
  rd 1:4268  
  export map my_export_map  
  route-target import 1:5230  
  route-target export 1:5231  
!  
interface loopback 0  
  ip address 10.1.2.1 255.255.255.255  
!  
interface atm 3/1.5 point-to-point  
  atm pvc 5 1 5000 aal5snap 0 0 0  
  mpls  
  mpls ldp profile default  
  description To P  
  ip address 172.16.1.2 255.255.255.252  
!  
  
router bgp 1  
  no synchronization  
  neighbor 10.1.2.0 remote-as 1  
  neighbor 10.1.2.0 update-source 10.1.2.1  
  neighbor 10.1.2.2 remote-as 1  
  neighbor 10.1.2.2 update-source 10.1.2.1  
  neighbor 10.1.2.3 remote-as 1
```

```
neighbor 10.1.2.3 update-source 10.1.2.1
!
address-family vpnv4 unicast
neighbor 10.1.2.0 activate
neighbor 10.1.2.0 next-hop-self
neighbor 10.1.2.0 send-community
neighbor 10.1.2.0 send-community extended
neighbor 10.1.2.2 activate
neighbor 10.1.2.2 next-hop-self
neighbor 10.1.2.2 send-community
neighbor 10.1.2.2 send-community extended
neighbor 10.1.2.3 activate
neighbor 10.1.2.3 next-hop-self
neighbor 10.1.2.3 send-community
neighbor 10.1.2.3 send-community extended
exit-address-family
!
address-family ipv4 vrf Orch_1_4268
no synchronization
no auto-summary
redistribute static
neighbor 192.168.0.2 remote-as 100
neighbor 192.168.0.2 update-source 192.168.0.1
neighbor 192.168.0.2 description s1.1
neighbor 192.168.0.2 send-community
neighbor 192.168.0.2 as-override
neighbor 192.168.0.2 maximum-prefix 10 warning-only
neighbor 192.168.0.2 prefix-list my_prefix_list_in in
neighbor 192.168.0.2 prefix-list my_prefix_list_out out
neighbor 192.168.0.2 allowas_in 4
neighbor 192.168.0.2 password v6lne0qkel33&
exit-address-family
-----
virtual-router PE3:Orch_1_4268
```

```
interface atm 3/0.6 point-to-point
  atm pvc 6 1 6000 aal5snap 0 0 0
  ip address 192.168.0.1 255.255.255.252
```


Index

Numerics

3-spoke hub site
 sample configuration 65

A

Allow AS in 44
Alternative routes, maintain in routing
 table 46
Applying configuration 17
AS override 44
ASN
 allow in AS_PATH 33
 override in AS_PATH 33, 44
Audit logs 56
Authentication 17
Authentication and communication 14

B

BGP
 autonomous systems 32
 communities 38
 domain-level parameters 32
 eBGP configuration 42

C

Capabilities 16
CE routers
 manual configuration 23
 roles 34
Check and force consistency 18
clear ip bgp command 61
clear ip route command 61
Co-existence with previously configured
 iBGP 42
Command-line interface 14, 17
Command-line parameters
 details 8
 on Solaris 10

 on Windows 10

Commands

clear ip bgp 61
clear ip route 61
mpls 22
mpls ldp 22
show config 58
show config virtual-router 58
show ip bgp 60
show ip bgp neighbor 59
show ip bgp vpn4 60
show ip route 59
show ip route vrf 60
show ip vrf 60
show mpls forwarding 60
show mpls interface 60
show processes cpu 58
show version 58

Communication and authentication 14

Communication problems 57

Configuration

 3-spoke hub site example 65
 applying 17
 spoke site example 64

Configuration timeout period 9

Configuring

 MPLS VPNs manually 22
 SNMP 22

customer support vii

D

Default route
 distributing 48
Detecting manual configuration 19
Device capabilities 16
Device discovery 14, 34
Device logs 56
Device roles 34

Device, configuration example 63
Discovery and MPLS VPNs 34
documentation
 downloading viii
 Service Activator ix

E

eBGP 43
 ASN in AS_PATH 44
 load sharing 46
 local preference 45
 prefix limits 46
 route dampening 47
eBGP commands 43
E-series commands
 general configuration 58
 VPN configuration 58
E-series support
 devices 2
 interfaces 2
 MPLS VPN features 4
 operating systems 2
 Service Activator features 4
Export maps
 configuring 24
 pre-defined 24, 40
Extended community attributes 33, 42, 45

H

Hub and spoke VPNs 37

I

iBGP
 MD5 authentication 42
 peering 32
iBGP configuration 41
IGP
 on PE routers 22
Installation issues 8
Interface
 roles 34
 types 2

L

Load balancing 46
Load sharing 46
Local preference 43
Log files 56

Loopback interface 22, 23, 33

M

Management VPNs 37
Manual configuration 19
Manual configuration for MPLS VPNs
 VRF tables 19
Maximum paths, in VPNs 33
MD5 authentication
 PE to PE 33, 42
Mesh VPNs 37
mpls command 22
mpls ldp command 22
MPLS tunnels 22
MPLS VPN configuration
 mandatory 22
 optional 23
MPLS VPNs 22
 device roles 34
 discovery 34
 domain-level parameters 32
 example configurations 63
 hub and spoke route targets 38
 PE-CE configuration 43, 48, 51
 PE-PE peering 41
 pre-requisites 32
 route distinguishers 36
 route targets 37, 38
 VRF tables 34
MPLS VPNs, VRF tables
 manually pre-configured, handling 39
Multiple paths, in VPNs 46

O

Operating system versions 2

P

P routers
 manual configuration 22
Passwords 14, 57
PE routers
 manual configuration 22
PE-CE configuration
 eBGP 43
 local preference 43
 static routes 51
 static routing 51
PE-PE peering 41

- MPLS VPNs 41
- ported 39
- Pre-defined
 - prefix list filters 25
 - route targets 38
 - VRF tables 23
 - VRF tables, removing 23, 40
- Pre-defined export maps 24
- Prefix limits 46
- Prefix list filters
 - examples 27
 - pre-defined 25
- Prefix lists
 - creating 26
- products
 - downloading viii
- R**
- RD number
 - per VPN 39
- RD numbers 36
- Read community 22
- Role assignment 34
- Route dampening 47
- Route distinguishers 36
- Route redistribution
 - into eBGP 47
- Route reflectors 32, 42
- Route targets 37
 - in hub and spoke VPNs 38
 - user-defined 38
- Router pre-configuration 21, 32
- Routes, allow load-sharing for VPNs 46
- RT numbers 37
- S**
- Securing TCP connections 42
- Security 17
- Shareable VRF tables 39
- show config command 58
- show config virtual-router command 58
- show ip bgp command 60
- show ip bgp neighbor command 59
- show ip bgp vpnv4 command 60
- show ip route command 59
- show ip route vrf command 60
- show ip vrf command 60
- show mpls forwarding command 60
- show mpls interface command 60
- show processes cpu command 58
- show version command 58
- Site of origin 46
- SNMP 22
 - configuring 22
 - read community 57
- SOO 46
- Spoke site
 - sample configuration 64
- Standard community attributes 33, 42, 45
- Static routing 51
- support
 - customer vii
- Supported devices 2
- Supported interfaces 2
- Supported models and operating systems 1
- Supported operating systems 2
- T**
- Troubleshooting 55
- V**
- Virtual routers
 - before discovery 14
 - capabilities 17
 - configuration 14
 - display 16
 - explanation 3
 - moving 16
 - VRF tables 34
- VPN configuration 58
 - troubleshooting 58
- VPN preservation 19
- VPN topologies 37
- VPNs
 - fully meshed 37
 - hub and spoke 37
 - management 37
 - maximum paths 33
 - multiple paths 46
- VPNs See MPLS VPNs
- VRF reduction 39
- VRF re-use 39
- VRF table names 35
- VRF tables 34
 - manually pre-configured, handling 39
 - pre-defined 23

pre-defined, removing 40
prefix limit 46