**JUNOS® Software**

# Multiplay Solutions Guide

*Release 9.3*

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

Part Number: 530-027213-01, Revision 1

*JUNOS® Software  Multiplay Solutions Guide*
Release 9.3


Writing: Mark Barnard, Justine Kangas, Sarah Lesway-Ball, Brian Wesley Simmons
Editing: Ben Mann, Stella Hackell
Illustration: Nathaniel Woodward, Mark Barnard
Cover Design: Edmonds Design

Revision History
10 October 2008—Revision 1

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

**READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE.** BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at http://www.gnu.org/licenses/gpl.html, and a copy of the LGPL at http://www.gnu.org/licenses/lgpl.html.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

# Abbreviated Table of Contents

# Table of Contents

# List of Figures

# List of Tables

# About This Guide

This preface provides the following guidelines for using the *JUNOS® Software Multiplay Solutions Guide*:



## Objectives

This guide describes how you can deploy IP television (IPTV) and voice over IP (VoIP) services in your network.

**NOTE:** This guide documents Release 9.3 of the JUNOS software. For additional information about the JUNOS software—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net/.

## Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M-series, MX-series, T-series, EX-series, or J-series routing platform.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery

- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

## Supported Routing Platforms

For the features described in this manual, the JUNOS software currently supports the following routing platforms:

- M-series
- MX-series
- T-series

## Using the Indexes

This reference contains a standard index with topic entries.

## Documentation Conventions

Table 1 on page xx defines notice icons used in this guide.

**Table 1: Notice Icons**

| Icon | Meaning | Description |
|---|---|---|
|  | Informational note | Indicates important features or instructions. |
|  | Caution | Indicates a situation that might result in loss of data or hardware damage. |
|  | Warning | Alerts you to the risk of personal injury or death. |
|  | Laser warning | Alerts you to the risk of personal injury from a laser. |

Table 2 on page xxi defines the text and syntax conventions used in this guide.

**Table 2: Text and Syntax Conventions**

| Convention | Description | Examples |
|---|---|---|
| Bold text like this | Represents text that you type. | To enter configuration mode, type the **configure** command:<br><br>user@host> **configure** |
| `Fixed-width text like this` | Represents output that appears on the terminal screen. | `user@host> show chassis alarms`<br>`No alarms currently active` |
| *Italic text like this* | ■ Introduces important new terms.<br>■ Identifies book names.<br>■ Identifies RFC and Internet draft titles. | ■ A policy *term* is a named structure that defines match conditions and actions.<br>■ *JUNOS System Basics Configuration Guide*<br>■ RFC 1997, *BGP Communities Attribute* |
| *Italic text like this* | Represents variables (options for which you substitute a value) in commands or configuration statements. | Configure the machine's domain name:<br><br>[edit]<br>root@# **set system domain-name** *domain-name* |
| Plain text like this | Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components. | ■ To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level.<br>■ The console port is labeled **CONSOLE**. |
| < > (angle brackets) | Enclose optional keywords or variables. | **stub <default-metric** *metric***>;** |
| \| (pipe symbol) | Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity. | **broadcast \| multicast**<br><br>(*string1* \| *string2* \| *string3*) |
| # (pound sign) | Indicates a comment specified on the same line as the configuration statement to which it applies. | **rsvp { # Required for dynamic MPLS only** |
| [ ] (square brackets) | Enclose a variable for which you can substitute one or more values. | **community name members [** *community-ids* **]** |

**Table 2: Text and Syntax Conventions** *(continued)*

| Convention | Description | Examples |
|---|---|---|
| Indention and braces ( { } ) | Identify a level in the configuration hierarchy. | [edit]<br>routing-options {<br>  static {<br>  route default {<br>  nexthop *address*;<br>  retain;<br>  }<br>  }<br>} |
| ; (semicolon) | Identifies a leaf statement at a configuration hierarchy level. | |
| **J-Web GUI Conventions** | | |
| **Bold text like this** | Represents J-Web graphical user interface (GUI) items you click or select. | ■ In the Logical Interfaces box, select **All Interfaces**.<br>■ To cancel the configuration, click **Cancel**. |
| > (bold right angle bracket) | Separates levels in a hierarchy of J-Web selections. | In the configuration editor hierarchy, select **Protocols > Ospf**. |

## List of Technical Publications

Table 3 on page xxii lists the software and hardware guides and release notes for Juniper Networks J-series, M-series, MX-series, and T-series routing platforms and describes the contents of each document. Table 4 on page xxvii lists the books included in the *Network Operations Guide* series. Table 5 on page xxvii lists the manuals and release notes supporting JUNOS software with enhanced services. All documents are available at http://www.juniper.net/techpubs/.

Table 6 on page xxviii lists additional books on Juniper Networks solutions that you can order through your bookstore. A complete list of such books is available at http://www.juniper.net/books.

**Table 3: Technical Documentation for Supported Routing Platforms**

| Book | Description |
|---|---|
| **JUNOS Software for Supported Routing Platforms** | |
| *Access Privilege* | Explains how to configure access privileges in user classes by using permission flags and regular expressions. Lists the permission flags along with their associated command-line interface (CLI) operational mode commands and configuration statements. |
| *Class of Service* | Provides an overview of the class-of-service (CoS) functions of the JUNOS software and describes how to configure CoS features, including configuring multiple forwarding classes for transmitting packets, defining which packets are placed into each output queue, scheduling the transmission service level for each queue, and managing congestion through the random early detection (RED) algorithm. |

**Table 3: Technical Documentation for Supported Routing Platforms** *(continued)*

| Book | Description |
|---|---|
| *CLI User Guide* | Describes how to use the JUNOS command-line interface (CLI) to configure, monitor, and manage Juniper Networks routing platforms. This material was formerly covered in the *JUNOS System Basics Configuration Guide*. |
| *Feature Guide* | Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software. |
| *High Availability* | Provides an overview of hardware and software resources that ensure a high level of continuous routing platform operation and describes how to configure high availability (HA) features such as nonstop active routing (NSR) and graceful Routing Engine switchover (GRES). |
| *MPLS Applications* | Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols. |
| *Multicast Protocols* | Provides an overview of multicast concepts and describes how to configure multicast routing protocols. |
| *Multiplay Solutions* | Describes how you can deploy IPTV and voice over IP (VoIP) services in your network. |
| *MX-series Layer 2 Configuration Guide* | Provides an overview of the Layer 2 functions of the MX-series routers, including configuring bridging domains, MAC address and VLAN learning and forwarding, and spanning-tree protocols. It also details the routing instance types used by Layer 2 applications. All of this material was formerly covered in the *JUNOS Routing Protocols Configuration Guide*. |
| *MX-series Layer 2 Solutions Guide* | Describes common configuration scenarios for the Layer 2 features supported on the MX-series routers, including basic bridged VLANs with normalized VLAN tags, aggregated Ethernet links, bridge domains, Multiple Spanning Tree Protocol (MSTP), and integrated routing and bridging (IRB). |
| *Network Interfaces* | Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform. |
| *Network Management* | Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options. |
| *Policy Framework* | Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options. |
| *Protected System Domain* | Provides an overview of the JCS 1200 platform and the concept of Protected System Domains (PSDs). The JCS 1200 platform, which contains up to six redundant pairs of Routing Engines running JUNOS software, is connected to a T320 router or to a T640 or T1600 routing node. To configure a PSD, you assign any number of Flexible PIC concentrators (FPCs) in the T-series routing platform to a pair of Routing Engines on the JCS 1200 platform. Each PSD has the same capabilities and functionality as a physical router, with its own control plane, forwarding plane, and administration. |

**Table 3: Technical Documentation for Supported Routing Platforms** *(continued)*

| Book | Description |
|---|---|
| *Routing Protocols* | Provides an overview of routing concepts and describes how to configure routing instances, and unicast routing protocols. |
| *Secure Configuration Guide for Common Criteria and JUNOS-FIPS* | Provides an overview of secure Common Criteria and JUNOS-FIPS protocols for the JUNOS software and describes how to install and configure secure Common Criteria and JUNOS-FIPS on a routing platform. |
| *Services Interfaces* | Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the router. |
| *Software Installation and Upgrade Guide* | Describes the JUNOS software components and packaging and explains how to initially configure, reinstall, and upgrade the JUNOS system software. This material was formerly covered in the *JUNOS System Basics Configuration Guide*. |
| *Subscriber Access* | Provides an overview of the subscriber access features of the JUNOS software and describes how to configure subscriber access support on the router, including dynamic profiles, class of service, AAA, and access methods. |
| *System Basics* | Describes Juniper Networks routing platforms and explains how to configure basic system parameters, supported protocols and software processes, authentication, and a variety of utilities for managing your router on the network. |
| *VPNs* | Provides an overview and describes how to configure Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits. Provides configuration examples. |
| **JUNOS References** | |
| *Hierarchy and RFC Reference* | Describes the JUNOS configuration mode commands. Provides a hierarchy reference that displays each level of a configuration hierarchy, and includes all possible configuration statements that can be used at that level. This material was formerly covered in the *JUNOS System Basics Configuration Guide*. |
| *Interfaces Command Reference* | Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces. |
| *Routing Protocols and Policies Command Reference* | Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing policies and protocols, including firewall filters. |
| *System Basics and Services Command Reference* | Describes the JUNOS software operational mode commands you use to monitor and troubleshoot system basics, including commands for real-time monitoring and route (or path) tracing, system software management, and chassis management. Also describes commands for monitoring and troubleshooting services such as class of service (CoS), IP Security (IPsec), stateful firewalls, flow collection, and flow monitoring. |

**Table 3: Technical Documentation for Supported Routing Platforms** *(continued)*

| Book | Description |
|---|---|
| *System Log Messages Reference* | Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message. |
| **J-Web User Guide** | |
| *J-Web Interface User Guide* | Describes how to use the J-Web graphical user interface (GUI) to configure, monitor, and manage Juniper Networks routing platforms. |
| **JUNOS API and Scripting Documentation** | |
| *JUNOScript API Guide* | Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms. |
| *JUNOS XML API Configuration Reference* | Provides reference pages for the configuration tag elements in the JUNOS XML API. |
| *JUNOS XML API Operational Reference* | Provides reference pages for the operational tag elements in the JUNOS XML API. |
| *NETCONF API Guide* | Describes how to use the NETCONF API to monitor and configure Juniper Networks routing platforms. |
| *JUNOS Configuration and Diagnostic Automation Guide* | Describes how to use the commit script and self-diagnosis features of the JUNOS software. This guide explains how to enforce custom configuration rules defined in scripts, how to use commit script macros to provide simplified aliases for frequently used configuration statements, and how to configure diagnostic event policies. |
| **Hardware Documentation** | |
| *Hardware Guide* | Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide. |
| *PIC Guide* | Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide. |
| *DPC Guide* | Describes the Dense Port Concentrators (DPCs) for all MX-series routers. |
| **JUNOScope Documentation** | |
| *JUNOScope Software User Guide* | Describes the JUNOScope software graphical user interface (GUI), how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations. |
| **Advanced Insight Solutions (AIS) Documentation** | |
| *Advanced Insight Solutions Guide* | Describes the Advanced Insight Manager (AIM) application, which provides a gateway between JUNOS devices and Juniper Support Systems (JSS) for case management and intelligence updates. Explains how to run AI-Scripts on Juniper Networks devices. |

**Table 3: Technical Documentation for Supported Routing Platforms** *(continued)*

| Book | Description |
|---|---|
| **J-series Routing Platform Documentation** | |
| *Getting Started Guide* | Provides an overview, basic instructions, and specifications for J-series routing platforms. The guide explains how to prepare your site for installation, unpack and install the router and its components, install licenses, and establish basic connectivity. Use the *Getting Started Guide* for your router model. |
| *Basic LAN and WAN Access Configuration Guide* | Explains how to configure the interfaces on J-series Services Routers for basic IP routing with standard routing protocols, ISDN backup, and digital subscriber line (DSL) connections. |
| *Advanced WAN Access Configuration Guide* | Explains how to configure J-series Services Routers in virtual private networks (VPNs) and multicast networks, configure data link switching (DLSw) services, and apply routing techniques such as policies, stateless and stateful firewall filters, IP Security (IPsec) tunnels, and class-of-service (CoS) classification for safer, more efficient routing. |
| *Administration Guide* | Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers. |
| **Release Notes** | |
| *JUNOS Release Notes* | Summarize new features and known problems for a particular software release, provide corrections and updates to published JUNOS, JUNOScript, and NETCONF manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures. |
| *Hardware Release Notes* | Describe the available documentation for the routing platform and summarize known problems with the hardware and accompanying software. Each platform has its own release notes. |
| *JUNOScope Release Notes* | Contain corrections and updates to the published JUNOScope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures. |
| *AIS Release Notes* | Summarize AIS new features and guidelines, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide initial setup, upgrade, and downgrade procedures. |
| *AIS AI-Scripts Release Notes* | Summarize AI-Scripts new features, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide instructions for automatic and manual installation, including deleting and rolling back. |
| *J-series Services Router Release Notes* | Briefly describe Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions. |

**Table 4: JUNOS Software Network Operations Guides**

| Book | Description |
|---|---|
| *Baseline* | Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages. |
| *Interfaces* | Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms. |
| *MPLS* | Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network. |
| *MPLS Log Reference* | Describes MPLS status and error messages that appear in the output of the show mpls lsp extensive command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network. |
| *MPLS Fast Reroute* | Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing. |
| *Hardware* | Describes tasks for monitoring M-series and T-series routing platforms. |

To configure and operate a J-series Services Router running JUNOS software with enhanced services, you must also use the configuration statements and operational mode commands documented in JUNOS configuration guides and command references. To configure and operate a WX Integrated Services Module, you must also use WX documentation.

**Table 5: JUNOS Software with Enhanced Services Documentation**

| Book | Description |
|---|---|
| **All Platforms** | |
| *JUNOS Software Interfaces and Routing Configuration Guide* | Explains how to configure J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification. |
| *JUNOS Software Security Configuration Guide* | Explains how to configure and manage security services such as stateful firewall policies, IP Security (IPsec) virtual private networks (VPNs), firewall screens, Network Address Translation (NAT), Public Key Cryptography, and Application Layer Gateways (ALGs). |

**Table 5: JUNOS Software with Enhanced Services Documentation** *(continued)*

| Book | Description |
|---|---|
| *JUNOS Software Administration Guide* | Shows how to monitor J-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems. |
| *JUNOS Software CLI Reference* | Provides the complete JUNOS software with enhanced services configuration hierarchy and describes the configuration statements and operational mode commands not documented in the standard JUNOS manuals. |
| **J-series Only** | |
| *JUNOS Software with Enhanced Services Design and Implementation Guide* | Provides guidelines and examples for designing and implementing IPsec VPNs), firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services. |
| *JUNOS Software with Enhanced Services Quick Start* | Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity. |
| *JUNOS Software with Enhanced Services J-series Services Router Hardware Guide* | Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications. |
| *JUNOS Software with Enhanced Services Migration Guide* | Provides instructions for migrating an SSG device running ScreenOS software or a J-series Services Router running the JUNOS software to JUNOS software with enhanced services. |
| *WXC Integrated Services Module Installation and Configuration Guide* | Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration. |
| *JUNOS Software with Enhanced Services for J-series Services Router Release Notes* | Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on J-series Services Routers, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services. |

**Table 6: Additional Books Available Through http://www.juniper.net/books**

| Book | Description |
|---|---|
| *Interdomain Multicast Routing* | Provides background and in-depth analysis of multicast routing using Protocol Independent Multicast sparse mode (PIM SM) and Multicast Source Discovery Protocol (MSDP); details any-source and source-specific multicast delivery models; explores multiprotocol BGP (MBGP) and multicast IS-IS; explains Internet Gateway Management Protocol (IGMP) versions 1, 2, and 3; lists packet formats for IGMP, PIM, and MSDP; and provides a complete glossary of multicast terms. |

**Table 6: Additional Books Available Through http://www.juniper.net/books** *(continued)*

| Book | Description |
|---|---|
| *JUNOS Cookbook* | Provides detailed examples of common JUNOS software configuration tasks, such as basic router configuration and file management, security and access control, logging, routing policy, firewalls, routing protocols, MPLS, and VPNs. |
| *MPLS-Enabled Applications* | Provides an overview of Multiprotocol Label Switching (MPLS) applications (such as Layer 3 virtual private networks [VPNs], Layer 2 VPNs, virtual private LAN service [VPLS], and pseudowires), explains how to apply MPLS, examines the scaling requirements of equipment at different points in the network, and covers the following topics: point-to-multipoint label switched paths (LSPs), DiffServ-aware traffic engineering, class of service, interdomain traffic engineering, path computation, route target filtering, multicast support for Layer 3 VPNs, and management and troubleshooting of MPLS networks. |
| *OSPF and IS-IS: Choosing an IGP for Large-Scale Networks* | Explores the full range of characteristics and capabilities for the two major link-state routing protocols: Open Shortest Path First (OSPF) and IS-IS. Explains architecture, packet types, and addressing; demonstrates how to improve scalability; shows how to design large-scale networks for maximum security and reliability; details protocol extensions for MPLS-based traffic engineering, IPv6, and multitopology routing; and covers troubleshooting for OSPF and IS-IS networks. |
| *Routing Policy and Protocols for Multivendor IP Networks* | Provides a brief history of the Internet, explains IP addressing and routing (Routing Information Protocol [RIP], OSPF, IS-IS, and Border Gateway Protocol [BGP]), explores ISP peering and routing policies, and displays configurations for both Juniper Networks and other vendors' routers. |
| *The Complete IS-IS Protocol* | Provides the insight and practical solutions necessary to understand the IS-IS protocol and how it works by using a multivendor, real-world approach. |

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Document part number
- Page number
- Software release version (not required for *Network Operations Guides [NOGs])*

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at http://www.juniper.net/customers/support/downloads/710059.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

### Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool located at https://tools.juniper.net/SerialNumberEntitlementSearch/.

### Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/ .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at http://www.juniper.net/support/requesting-support.html.

# Part 1
# IPTV Network Solutions

Chapter 1

# IPTV Video Application

Next-generation multiplay networks are voice, data, and video networks that support personalized media and interactive IP television (IPTV) services along with communications services such as voice over IP (VoIP) and Internet data transmission. These services place extreme demands on network scalability, quality of service, security, and bandwidth resources. JUNOS software provides support for configuring various broadband video architectures in a multiplay network.

Although the overview in this chapter discusses more than one video network model, the example focuses on one specific video network architecture that incorporates a video services router running JUNOS software Release 8.3 or later. To understand this chapter, you should be familiar with Broadband Remote Access Server (B-RAS) operation on Juniper Networks routers, as well as standard Internet Group Management Protocol (IGMP) configurations.

For more information about B-RAS configuration, see the *JUNOSe Broadband Access Configuration Guide*. For more information about IGMP configuration, see the *JUNOS Multicast Protocols Configuration Guide* or *JUNOSe Multicast Routing Configuration Guide*. You can obtain these manuals at:

http://www.juniper.net/techpubs/software/index.html

This chapter covers the following topics:



## System Requirements

To implement video services on a routing platform running JUNOS software, you must use the following software and hardware components:

- JUNOS Release 8.3 or later for next-generation broadband or video features
- Juniper Networks video services routers (for example, the MX960 Ethernet Services Router or any M-series router that supports the JUNOS Release 8.3 or later video services software package)

## Terms and Acronyms

- **ASM (Any Source Multicast)**—A method of allowing a multicast receiver to listen to all traffic sent to a multicast group, regardless of its source.
- **BSR (broadband services router)**—A router used for subscriber management and edge routing.
- **IGMP (Internet Group Membership Protocol)**—A host to router signaling protocol for IPv4 used to support IP multicasting.
- **IS-IS (Intermediate System-to-Intermediate System)**—A link-state, interior gateway routing protocol (IGRP) for IP networks that uses the shortest-path-first (SPF) algorithm to determine routes.
- **LSP (label-switched path)**—The path traversed by a packet that is routed by MPLS. Some LSPs act as tunnels. LSPs are unidirectional, carrying traffic only in the downstream direction from an ingress node to an egress node.
- **MPLS (Multiprotocol Label Switching)**—A mechanism for engineering network traffic patterns that functions by assigning to network packets short labels that describe how to forward the packets through the network.
- **OIF (outgoing interface)**—An interface used by multicast functions within a router to determine which egress ports to use for fowarding multicast groups.
- **OSPF (Open Shortest Path First)**—A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).
- **PIM (Protocol Independent Multicast)**—A multicast routing protocol used for delivering multicast messages in a routed environment.
- **routing gateway**—A firewall, Network Address Translation (NAT) router, or other routing device used as a customer premises equipment (CPE) terminator in the home, office, or local point of presence (POP).
- **SSM (single-source multicast)**—A routing method that allows a multicast receiver to detect only a specifically identified sender within a multicast group.
- **set-top box**—The end host or device used to receive IPTV video streams.
- **VOD (video on demand)**—A unicast streaming video offering by service providers that enables the reception of an isolated video session per user with rewind, pause, and similar VCR-like capabilities.
- **VSR (video services router)**—A router used in a video services network to route video streams between an access network and a metro or core network. The VSR is any M-series or MX-series router that supports the video routing package provided with JUNOS software Release 8.3 or later.

## Overview and Topology

As an emerging genre of service, Internet Protocol television (IPTV) networks compete with more traditional video service offerings. IPTV networks provide new revenue streams to higher-premium multiplay services, which encompass bundled voice, video, Internet, gaming, and other services.

IPTV offers true integration of information, communications, and entertainment into personalized and interactive applications centered on familiar television-like services, including:

- Interactive entertainment services
- Broadcast services in standard and high-definition formats
- Video on demand (VOD)
- Digital video recording (DVR), including pausing and recording of broadcast TV, rewind, and fast-forward functionality
- Enhanced user services or interfaces such as an interactive programming guide and Mosaic interface, and converged features such as caller ID and message waiting

These new opportunities also present challenges to cost-effectively manage the delivery of performance-sensitive services over a service provider's IP infrastructure. Ensuring quality of service (QoS) for IPTV is essential, especially when the network is also carrying a wide array of other traffic. IPTV and similar latency-sensitive and jitter-sensitive services cannot be delivered at an acceptable quality of service simply through additional bandwidth. IPTV services must provide more efficient resource utilization while offering the best level of experience possible for subscribers.

Figure 1 on page 6 shows a basic video network topology. The example in this chapter uses this topology. This network topology can be viewed as having two parts: an access side and a metro/core side. The demarcation of these two parts is at the video services router.

**Figure 1: Basic Video Network Topology**

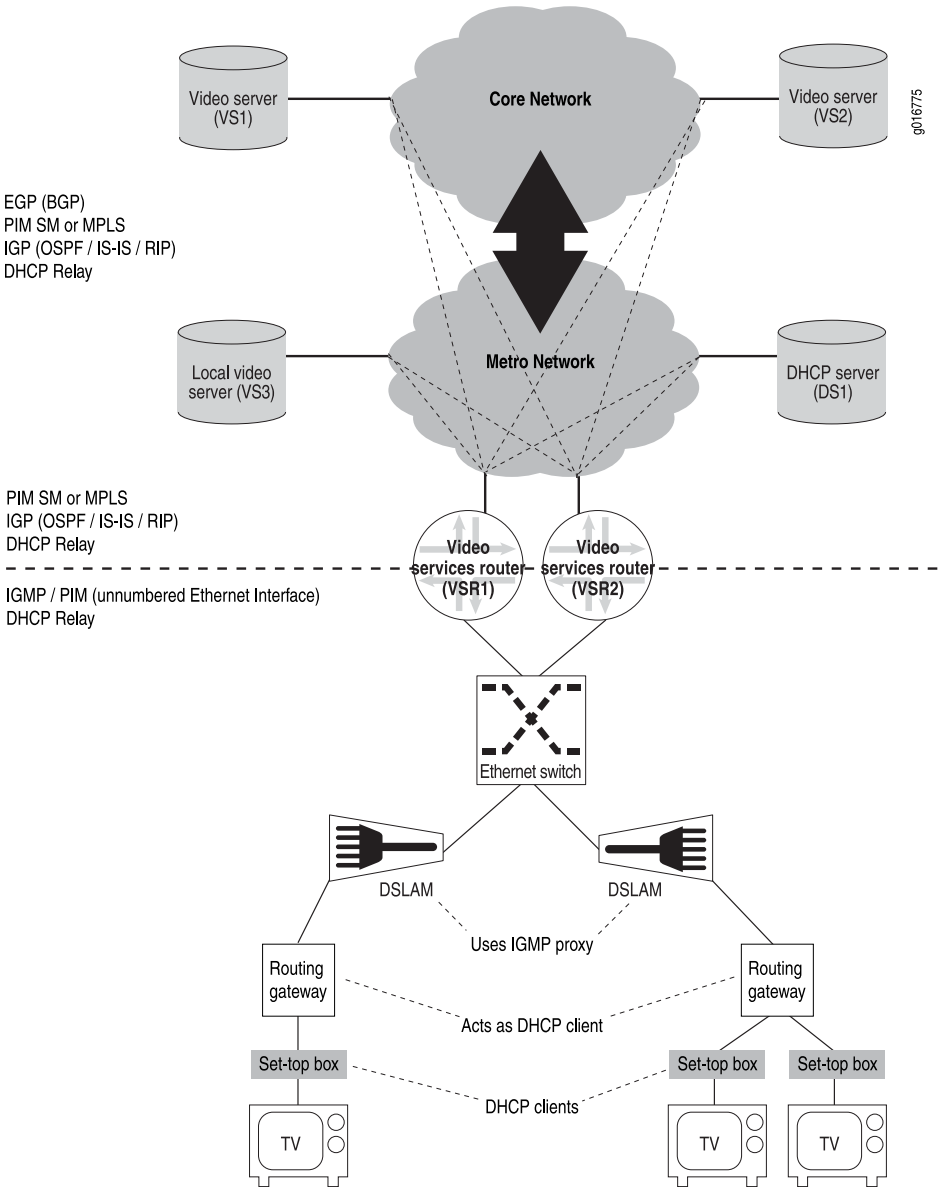


## *Video Network Elements*

The basic video (IPTV) network model, shown in Figure 2 on page 7, consists of up to five network elements.

**Figure 2: Basic IPTV Network Model**

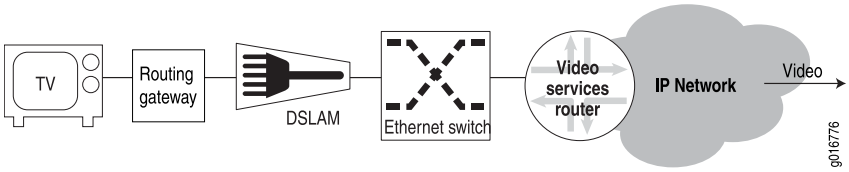


These network elements are:

- Set-top box

  At the subscriber site, a set-top box links the television to the external network. This device initiates channel change requests and responds to status inquiries.

- Routing gateway

  The routing gateway, often close to the subscriber site or a part of the set-top box, aggregates traffic from multiple subscribers and may act upon requests from the set-top box.

- DSLAM

  The digital subscriber line access multiplier (DSLAM), like the routing gateway, aggregates traffic from multiple subscribers and may act on requests from the set-top box. The DSLAM often resides at a separate, centrally located office.

- Ethernet switch

  Some networks can include an Ethernet switch or some other broadband services aggregator (BSA) to provide an additional layer of aggregation.

- Edge router

  The edge router (typically a broadband services router [BSR] or video services router [VSR]) is the gateway into the backbone network. This device most often controls the multicast traffic to and channel requests from the set-top box.

## *IGMP and Video Networks*

In a video (IPTV) network, broadcast television, pay-per-view (PPV), and video-on-demand (VOD) channels are all delivered by means of IP multicasting. Internet Group Management Protocol (IGMP) is the mechanism that controls the delivery of multicast traffic to subscribers on the network. This traffic is received and controlled by the subscriber's set-top box through multicast streams (referred to as channels). IGMP communicates with the upstream routing equipment to begin sending (join) or stop sending (leave) a channel.

Depending on the architecture that you choose for your network, the process of controlling channels occurs on a DSLAM, an aggregation switch, or an edge router.

## IGMP Basics

Basic IGMP operation involves the following two devices:

- IGMP host (client)—Device that issues messages to join or leave a multicast group. This device also responds to queries from the multicast router. A set-top box is an example of an IGMP host.
- IGMP router (multicast router)—Device that responds to the join and leave messages to determine whether or not to forward multicast groups from an interface. Periodic queries assist the router in recovering from any error conditions and verifying requests. The IGMP router receives multicast groups through the use of a multicast protocol, such as Protocol Independent Multicast (PIM), or through static flooding. An IGMP router is the termination point for any IGMP messages and therefore does not send any IGMP information to its upstream neighbors.

The IGMP protocol provides the following three basic functions for IP multicast networks:

- Join messages—Messages that indicate an IGMP host wants to receive information from (that is, become a member of) a multicast group.
- Leave messages—Messages that indicate an IGMP host no longer wants to receive information from a multicast group.
- Query messages—Messages from an IGMP router requesting information from a host. For example, if a set-top box is unplugged without first issuing a leave message, the IGMP router may query the host to determine what multicast groups the host belongs to.

## IGMP and Intermediate Devices

In early IGMP networks, devices located between the IGMP client and the IGMP router did not detect IGMP flows. In Figure 3 on page 8, the top set-top box issues a request to view Channel 4, and the DSLAM forwards the request to the edge router. In response, the edge router begins forwarding the multicast group associated with Channel 4. However, if it does not detect IGMP flows, the intermediate device (in this case, the DSLAM) cannot appropriately forward the multicast traffic. By default, most switches broadcast incoming multicast traffic to all ports. In this case, the broadcast results in the bottom client receiving an unrequested channel.

**Figure 3: DSLAM Without IGMP Flow Recognition**

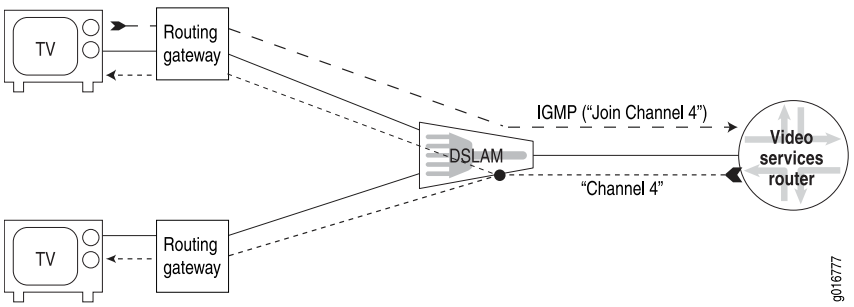

In these early networks, broadcasting of unrequested channels was not considered a problem, because multicast usage was low and the intermediate devices were typically LAN switches with lower interface and bandwidth costs. Now that IPTV requires higher bandwidth (often 4 Mbps per channel) and bandwidth costs more, it is crucial to ensure that IPTV channels are forwarded only to those subscribers currently viewing them.

To provide more intelligent control of bandwidth, DSLAMs and other intermediate devices now recognize IGMP flows. These devices examine incoming flows and build outgoing interface (OIF) tables. Figure 4 on page 9 shows a simple example of an outgoing interface table for the DSLAM. The outgoing interface table enables the DSLAM to appropriately forward each multicast group (or channel) from the correct port.

**Figure 4: DSLAM with IGMP Flow Recognition**

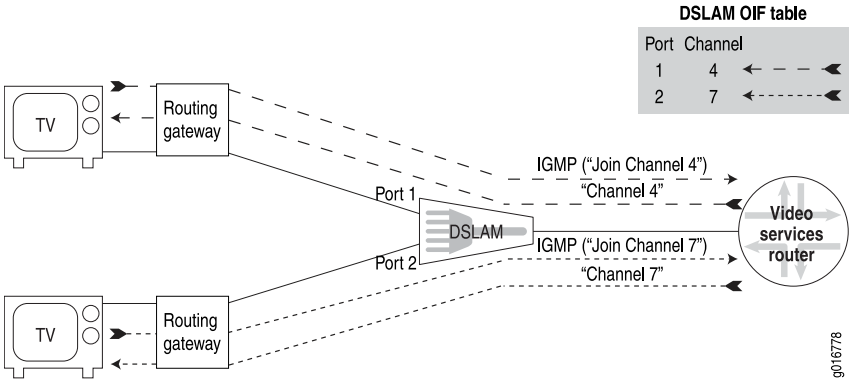


The intermediate device builds the outgoing interface table in one of two ways—IGMP snooping or IGMP proxy.

**CAUTION:** Some intermediate devices implement IGMP subsystems that use characteristics of both IGMP snooping and IGMP proxy. Most commonly, these devices might determine whether to forward IGMP packets (IGMP proxy) but do not modify the source IP address (IGMP snooping). We recommend that you avoid these nonstandard implementations.

## IGMP Snooping

Figure 5 on page 10 illustrates IGMP snooping, in which an intermediate device (such as a DSLAM) transparently monitors IGMP traffic. The device adds interfaces to its outgoing interface table when it detects join request messages and removes interfaces from its outgoing interface table when it detects leave request messages. The snooping device also maintains state information for general *membership query maximum response time* timers if the IGMP client does not issue a leave message (for example, if an IPTV set-top box experiences a power outage).

**Figure 5: IGMP Snooping**

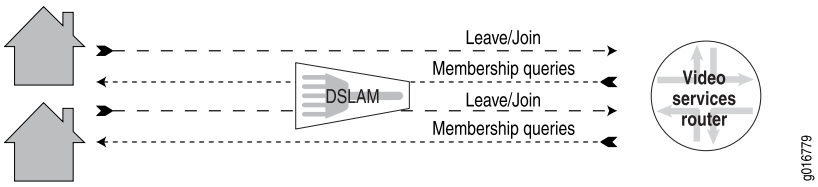


Because IGMP snooping is transparent, the snooping device typically does not participate in IGMP host messaging. The device only monitors transactions between clients and routers, forwarding IGMP packets upstream to the multicast router and determining when join or leave processing is required for a downstream host. One exception to this transparency occurs when the snooping device intercepts membership reports based on local filters to prevent the host from joining specific groups (that is, specific broadcast channels allocated to multicast groups that are blocked from being received by the set-top box).

The snooping device can receive multicast data in several ways within a broadband access network. The router might be configured to flood all multicast groups downstream to the snooping device. The upstream router might forward only groups based on IGMP membership reports that it receives from the IGMP hosts. The snooping agent might invoke an IGMP client process to source its own membership reports that it sends to the multicast router, and so on. However, these various options are beyond the scope of this document.

## IGMP Proxy

Figure 6 on page 10 illustrates an IGMP proxy. An IGMP proxy performs functions of both an IGMP router and an IGMP client. When an IGMP host issues a join message, the IGMP proxy receives the message and adds the interface to its outgoing interface table for a specific multicast group. The proxy uses a general membership query timer and state to send general queries downstream to all multicast-enabled interfaces. When the IGMP proxy receives a leave message, the proxy issues a group-specific query. If no hosts respond to the query within a configured response time interval, the proxy removes the interface from the outgoing interface table.

**Figure 6: IGMP Proxy**

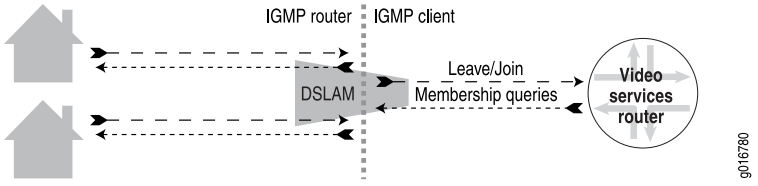


A device that functions as an IGMP proxy participates in every IGMP flow. This level of participation requires much more processing power and memory allocation from the DSLAM, but it can save upstream bandwidth.

Because a multicast router treats any IGMP proxy that it interacts with as an IGMP client, the multicast router tracks one device (the DSLAM) joining and leaving multicast

groups. As a result, the multicast router receives no information regarding subscribers on the other side of the IGMP proxy.

## DHCP Relay and Video Services Routers

The Dynamic Host Configuration Protocol (DHCP) provides an automated mechanism for network devices to obtain configuration information and a lease for an IP address.

The most important configuration parameter carried by DHCP is the IP address. A computer must initially be assigned a specific IP address that is appropriate to the network to which the computer is attached and that is not assigned to any other computer on that network. If you move a computer to a new network, it must be assigned a new IP address for that new network. You can use DHCP to manage these assignments automatically.

DHCP carries other important configuration parameters, such as the subnet mask, default router, and DNS server.

The video services router must run DHCP relay to enable devices to obtain parameters from the DHCP server on the network. The DHCP relay feature relays a request from a remote client to a DHCP server for an IP address. When the router receives a DHCP request from an IP client, it forwards the request to the DHCP server and passes the response back to the IP client.

For more information about configuring DHCP relay, see the *JUNOS Policy Framework Configuration Guide*.

## Video Networking and the Metro or Core Network

Video networks can incorporate various protocols used in the metro and core network. How you configure a metro or core network to transmit video streams depends on the type of network you have and the complexity of your application. All the protocols create multicast trees over which video streams can travel from one (or many) sources to a number of hosts.

### What IP Routing Protocols to Use

When running video networks in an IP metro and core network, you must configure several protocols to function together. These protocols typically include the following protocol types:

- A multicast protocol to route multicast traffic
- An interior gateway protocol (IGP) to provide topological information to the multicast protocol
- An exterior gateway protocol (EGP) to route between different networks (depending on the complexity of your network)

**Multicast Protocols to Use**—Video networks often use Protocol Independent Multicast sparse mode (PIM SM) when communicating beyond the access side of the network (that is, in the metro or core networks).

PIM is a family of multicast routing protocols that enable one-to-many and many-to-many distribution of data. The term *protocol-independent* means that PIM is not dependent on any particular unicast routing protocol for topology discovery. However, because it does not have its own method of topology discovery, PIM obtains routing information (such as dynamic endpoints) from other routing protocols, such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS).

Instead of flooding packets throughout the network and then removing (or pruning) paths where no receivers exist, PIM SM uses the information it receives from the other routing protocols to construct a tree from each sender to the receivers in a multicast group.

**Interior Gateway Protocols to Use**—PIM must use an IGP to obtain current topology information. The two protocols most often used by PIM to obtain topology information are OSPF and IS-IS. As IGPs, OSPF and IS-IS function within a single autonomous system (OSPF) or area (IS-IS).

Both OSPF and IS-IS are link-state routing protocols; they flood topology information throughout a network of routers within the autonomous system or area. After obtaining this information, each router independently builds a picture of the network topology. The routers can then forward packets or datagrams based on the best topological path through the network to the destination.

**Exterior Gateway Protocols to Use**—Depending on the complexity and size of your network, you might need to configure an exterior gateway protocol (EGP). EGPs such as Border Gateway Protocol (BGP) exchange routing information between networks.

### Using MPLS and Label-Switched Paths

Instead of using PIM SM to create multicast trees, you can use MPLS to control the paths that traffic takes to various destinations.

In the traditional Layer 3 forwarding paradigm, as a packet travels from one router to the next, an independent forwarding decision is made at each hop. Each device analyzes the IP network layer header and then chooses the next hop based on the analysis and the information in the routing table.

In an MPLS environment, however, the packet header is analyzed only once, when the packet enters the MPLS network. After analyzing the packet header, the router assigns the packet to a stream that is identified by a label (a short, fixed-length value at the front of the packet). Downstream routers use these labels as lookup indexes for the label-forwarding table. The label-forwarding table stores forwarding information for each label.

A point-to-multipoint MPLS label-switched path (LSP) is an RSVP-signaled LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

Point-to-multipoint LSPs enable you to do the following:

- Use MPLS for point-to-multipoint data distribution similar to that provided by IP multicast.
- Add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- Configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- Use link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be switched quickly to the bypass.
- Configure subpaths either statically or dynamically.
- Specify graceful restart on point-to-multipoint LSPs.

For additional information about how to configure MPLS point-to-multipoint LSPs, see the *JUNOS MPLS Applications Configuration Guide*.

### *Redundancy and Failure Detection for Video Services Routers*

Video networks require rapid failure detection and router redundancy to ensure minimal interruption of service. To provide a high level of failure detection and redundancy, you can employ PIM Bidirectional Forwarding Detection (BFD) for multicast traffic and Virtual Router Redundancy Protocol (VRRP) for unicast traffic in your video network.

## Sample Configuration of an IPTV Network

This section provides a comprehensive sample configuration for the video services routers (VSR1 and VSR2) in the network topology shown in Figure 1 on page 6 and described in the following example sections.

**Configuration for Router VSR1**

```
[edit]
interfaces {
   ge-1/0/0 {
      unit 0 {
         family inet {
            unnumbered-address lo0.0;
         }
      }
   }
   ge-1/0/1 {
      unit 0 {
         family inet {
            address 10.1.1.1/24;
         }
      }
      vlan-tagging;
      unit 1 {
```

```
            family inet {
               address 10.1.1.1/24 {
                  vrrp-group 1 {
                     virtual-address 10.1.1.99;
                     priority 200;
                     fast-interval 250;
                  }
               }
            }
         }
      }
   }
   protocols {
      igmp {
         interface ge-1/0/0.0;
         promiscuous-mode;
         immediate-leave;
      }
   }
   ospf {
      area 0 {
         interface ge-1/0/1;
      }
   }
   pim {
      rp {
         local {
            address 1.1.1.1;
         }
      }
      interface ge-1/0/0.0 {
         mode sparse;
         bfd-liveness-detection {
            minimum-interval 100;
         }
      }
      rp {
         local {
            address 1.1.1.1;
         }
      }
      interface ge-1/0/1.0 {
         mode sparse;
         bfd-liveness-detection {
            minimum-interval 100;
         }
      }
   }
   forwarding-options {
      dhcp-relay {
         server-group {
            DS1 {
               100.1.1.1;
            }
         }
         active-server-group DS1;
```

```
      group one {
        interface ge-1/0/0.0;
      }
    }
  }
  routing-options {
    static {
      route 1.1.1.1/32 {
        qualified-next-hop ge-1/0/0.0;
      }
    }
  }
```

**Configuration for Router VSR2**

```
[edit]
interfaces {
  ge-1/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
  ge-1/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.2/24;
      }
    }
  }
  ge-1/0/1 {
    vlan-tagging;
    unit 1 {
      family inet {
        address 10.1.1.2/24 {
          vrrp-group 1 {
            virtual-address 10.1.1.99;
            priority 100;
            fast-interval 250;
          }
        }
      }
    }
  }
}
protocols {
  igmp {
    interface ge-1/0/0.0;
    promiscuous-mode;
    immediate-leave;
  }
  ospf {
    area 0 {
      interface ge-1/0/1;
```

```
        }
    }
    pim {
        rp {
            local {
                address 1.1.1.1;
            }
        }
        interface ge-1/0/0.0 {
            mode sparse;
            bfd-liveness-detection {
                minimum-interval 100;
            }
        }
        rp {
            local {
                address 1.1.1.1;
            }
        }
        interface ge-1/0/1.0 {
            mode sparse;
            bfd-liveness-detection {
                minimum-interval 100;
            }
        }
    }
}
forwarding-options {
    dhcp-relay {
        server-group {
            DS1 {
                100.1.1.1;
            }
        }
        active-server-group DS1;
        group one {
            interface ge-1/0/0.0;
        }
    }
}
routing-options {
    static {
        route 1.1.1.2/32 {
            qualified-next-hop ge-1/0/0.0;
        }
    }
}
```

## Configuring the Access Side of a Video Services Router Running JUNOS Software

The access (or customer) side of the router running JUNOS software and operating in a video network uses IGMP and DHCP to manage video traffic to various clients. The interfaces on this side of the network use an unnumbered Ethernet configuration, as shown in Figure 7 on page 17.

**Figure 7: IPTV Network (Access Side)**

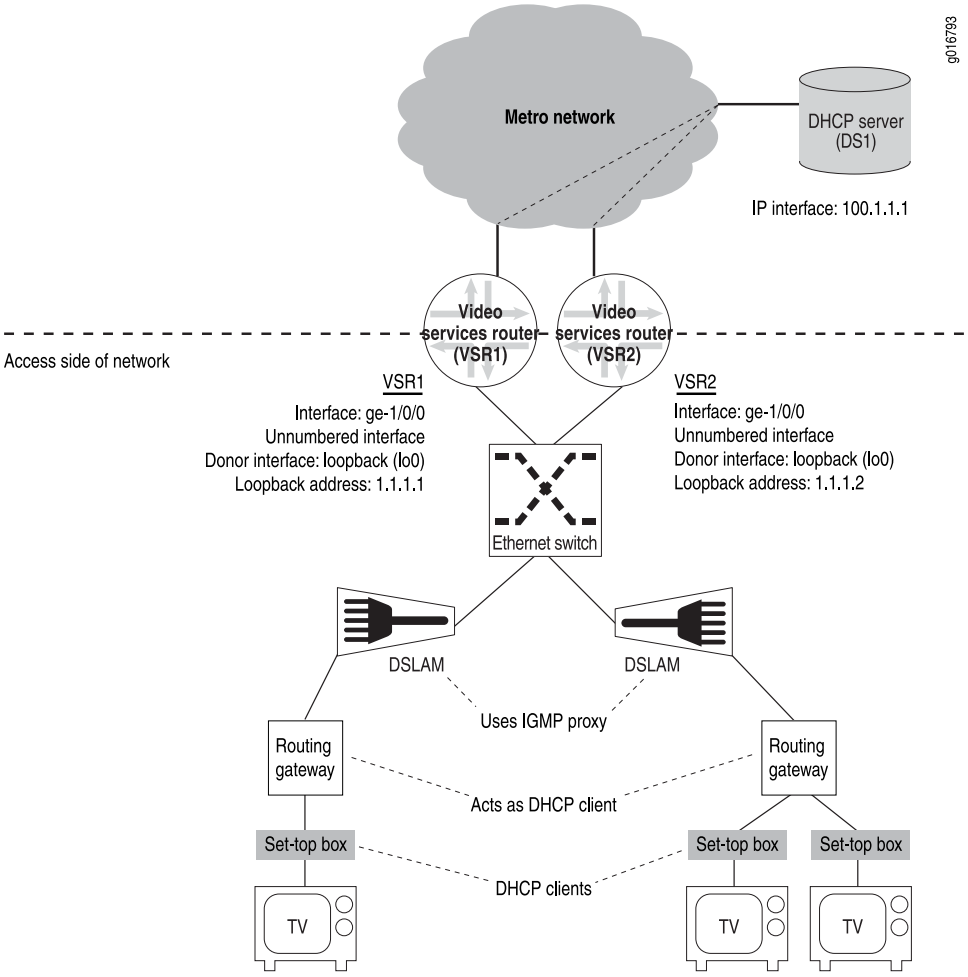


To implement video/IPTV applications on the access side of a video services router running JUNOS software, use the following procedures.

**NOTE:** To simplify this example, both video services routers (VSR1 and VSR2) use the same configuration except where otherwise specified.

1. Configure each access interface as an unnumbered Ethernet interface.

```
[edit]
interfaces {
   ge-1/0/0 {
      unit 0 {
         family inet {
            unnumbered-address lo0.0;
         }
      }
   }
```

```
}
```

2. Specify that the interface use promiscuous mode.

**NOTE:** You must specify that the IGMP interface use promiscuous mode if you define the unnumbered Ethernet donor interface as a loopback interface.

3. Specify that the IGMP interface use immediate leave if you want the interface to do one of the following:

   - For IGMPv2: Immediately remove the group membership from the interface and suppress the sending of any group-specific queries for the multicast group.
   - For IGMPv3: Suppress the sending of group-and-source queries and rely on the JUNOS-supported host tracking mechanism to determine group membership removal.

```
[edit]
protocols {
   igmp {
      interface ge-1/0/0.0;
      promiscuous-mode;
      immediate-leave;
   }
}
```

4. Configure DHCP relay.

```
[edit]
forwarding-options {
   dhcp-relay {
      server-group {
         DS1 {
            100.1.1.1; # IP address of DHCP server (DS1)
         }
      }
      active-server-group DS1;
      group one {
         interface ge-1/0/0.0; # interface to which DHCP clients send requests
      }
   }
}
```

5. Configure PIM (required to configure PIM BFD).

```
[edit]
protocols {
   pim {
      rp {
         local {
            address 1.1.1.1;
         }
      }
```

```
        interface ge-1/0/0.0 {
            mode sparse;
        }
    }
}
```

6. Configure PIM BFD to enable rapid failover detection for the PIM interfaces.

```
[edit]
protocols {
    pim {
        interface ge-1/0/0.0 {
            bfd-liveness-detection {
                minimum-interval 100;
            }
        }
    }
}
```

7. Configure static routes over which each loopback interface can communicate with the other.

   a. Configure a static route on Router VSR1 to the loopback interface on Router VSR2.

```
[edit]
routing-options {
    static {
        route 1.1.1.2/32 {
            qualified-next-hop ge-1/0/0.0;
        }
    }
}
```

   b. Configure a static route on Router VSR2 to the loopback interface on Router VSR1.

```
[edit]
routing-options {
    static {
        route 1.1.1.1/32 {
            qualified-next-hop ge-1/0/0.0;
        }
    }
}
```

**NOTE:** You must also configure static routes for the DSLAM devices to communicate with the unnumbered interfaces that are using a loopback interface.

## Configuring the Metro and Core Side of a Video Services Router Running JUNOS Software

The metro and core side of a router running JUNOS software and operating in a video network uses PIM SM or MPLS point-to-multipoint LSPs to manage video flows from various servers.

**Figure 8: IPTV Network (Metro and Core Side)**

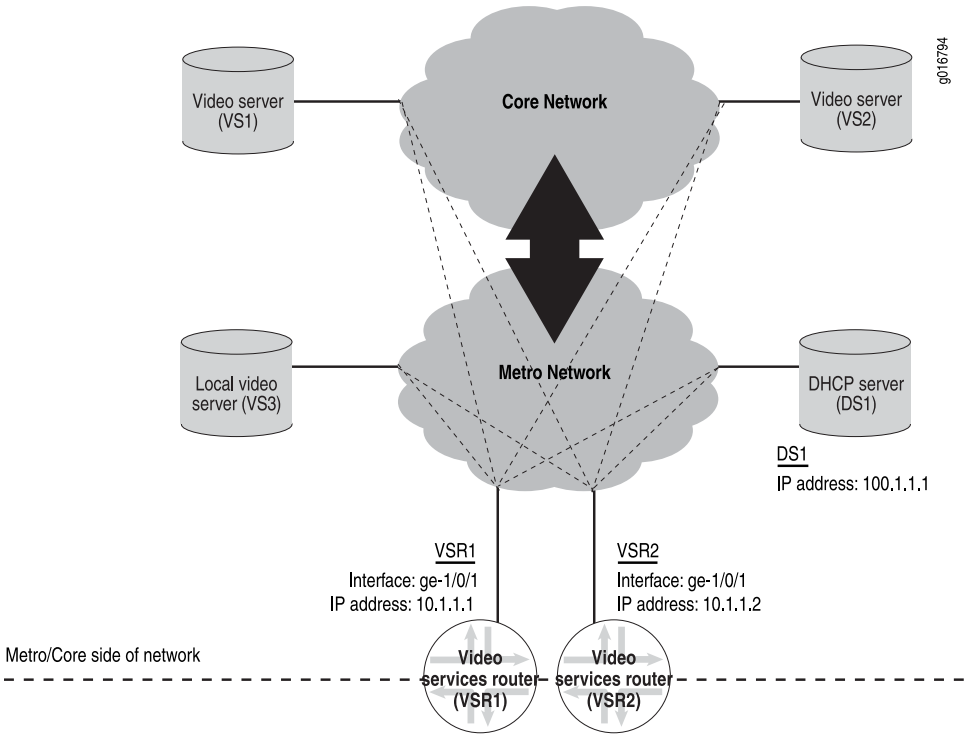


When using PIM SM, you must also configure an Internal Gateway Protocol (IGP) to dynamically maintain a topology of the network that PIM SM can use.

To implement video applications on the metro and core side of a video services router running JUNOS software, use the following procedures:

1. Configure static IP addresses for the metro and core interface for both Router VSR1 and VSR2.

   a. Configure a static IP address for Router VSR1.

```
[edit]
interfaces {
   ge-1/0/1 {
      unit 0 {
         family inet {
            address 10.1.1.1/24;
         }
      }
   }
```

```
}
```

b. Configure a static IP address for Router VSR2.

```
[edit]
interfaces {
   ge-1/0/1 {
      unit 0 {
         family inet {
            address 10.1.1.2/24;
         }
      }
   }
}
```

**NOTE:** You must define IP connectivity from the DHCP server (DS1) to the metro and core interface of Router VSR1 and VSR2.

2. Configure an internal gateway protocol (IGP). This example uses OSPF as the IGP for the network.

```
[edit]
protocols {
   ospf {
      area 0 {
         interface ge-1/0/1;
      }
   }
}
```

3. Configure PIM sparse mode (PIM SM).

**NOTE:** By default, IGMP is automatically enabled on all interfaces on which you configure PIM.

```
[edit]
protocols {
   pim {
      rp {
         local {
            address 1.1.1.1; # IP address of the PIM rendezvous point router
         }
      }
      interface ge-1/0/1.0 {
         mode sparse; # Define PIM SM on the metro and core interface
      }
   }
}
```

4. Configure PIM BFD to enable rapid failover detection for the PIM interfaces.

```
[edit]
protocols {
   pim {
      interface ge-1/0/1.0 {
         bfd-liveness-detection {
            minimum-interval 100;
         }
      }
   }
}
```

## Configuring Router Redundancy

The bidirectional forwarding detection (BFD) protocol that you configured on each PIM interface uses control packets and shorter detection time limits to detect failures rapidly in a network for multicast traffic. However, to configure redundancy for unicast traffic in a video network (for example, for video-on-demand streams), you can use Virtual Router Redundancy Protocol (VRRP).

VRRP enables hosts on a LAN to use redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts.

At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, thus always providing a virtual default routing platform and allowing traffic on the LAN to be routed without relying on a single routing platform.

To configure VRRP for each metro and core interface on the video services router, follow these steps.

**NOTE:** The VRRP groups must be the same for each router, and the VRRP priority setting must be lower for one of the routers.

1. Include the vrrp-group statement on the metro and core interface of Router VSR1.

```
[edit]
interfaces {
   ge-1/0/1 {
      vlan-tagging;
      unit 1 {
         family inet {
            address 10.1.1.1/24 {
               vrrp-group 1 {
                  virtual-address 10.1.1.99;
                  priority 200;
                  fast-interval 250;
               }
            }
         }
```

```
        }
    }
}
```

2. Include the vrrp-group statement on the metro and core interface of Router VSR2.

```
[edit]
interfaces {
    ge-1/0/1 {
        vlan-tagging;
        unit 1 {
            family inet {
                address 10.1.1.2/24 {
                    vrrp-group 1 {
                        virtual-address 10.1.1.99;
                        priority 100;
                        fast-interval 250;
                    }
                }
            }
        }
    }
}
```

## Verifying Your Configuration

You can use several commands to verify that the IPTV network is functioning and to monitor its status.

### *Verifying Connectivity*

When you configure your IPTV network, we recommend that you verify connectivity between routers (VSR1 and VSR2) and between each router and certain devices using the ping command.

The format for the ping command is as follows:

ping *host* source *source-address*

*host*—IP address of the device or interface to which you want to issue the ping command.

source *source-address*—IP address of the outgoing interface.

To verify connectivity:

- Issue the ping command from the access interface on each router to the loopback interface of the redundant router.
- Issue the ping command from each metro and core interface on each router to the metro and core interface on the redundant router.
- Issue the ping command from the loopback interface of each router to the DHCP server.

## *Using Operational Commands*

You can use various operational commands to obtain information about the IPTV network and to verify that the network is operating properly. Table 7 on page 24 lists specific operational commands that can provide information about the IPTV network and the protocols that you configured on each video services router.

**Table 7: Operational Commands for Network Verification**

| Operational Command | Purpose |
|---|---|
| show dhcp relay binding | The expected DHCP address bindings appear in the Dynamic Host Configuration Protocol (DHCP) client table. |
| show dhcp relay statistics | DHCP relay statistics are in line with expectations. |
| show igmp group | IGMP group membership is functioning as expected. |
| show igmp interface | ■ The status of each configured IGMP interface is operational (up).<br>■ The expected number of groups appears on each IGMP interface.<br>■ Promiscuous mode is enabled (on) for IGMP unnumbered interfaces. |
| show pim interfaces | ■ The status of each PIM interface is operational (up).<br>■ Each interface is running sparse mode. |
| show pim join | ■ PIM group joins are occurring as expected.<br>■ Each join is receiving sparse mode entries. |
| show pim neighbors<br><br>show pim neighbors detail | ■ PIM is establishing neighbor adjacencies correctly.<br>■ BFD is enabled. |
| show pim rps | The PIM rendezvous point router is correct. |
| show pim statistics | PIM statistics are in line with expectations. |

For additional information about these operational mode commands, see the *JUNOS Routing Protocols and Policies Command Reference*.

**Related Topics** Because the concepts that constitute logical routers cut across the entire JUNOS software documentation set, the following manuals can be useful references:

- For additional information about routing protocols, see the *JUNOS Routing Protocols Configuration Guide*
- For additional information about interface configuration, see the *JUNOS Network Interfaces Configuration Guide*

- For additional information about MPLS and related protocols, see the *JUNOS MPLS Applications Configuration Guide*
- For additional information about multicast protocols, configuring flow maps and flow cache properties, and configuring bandwidth management, see the *JUNOS Multicast Protocols Configuration Guide*
- For additional information about operational mode commands and output, see the *JUNOS Interfaces Command Reference*, the *JUNOS Routing Protocols and Policies Command Reference*, and the *JUNOS System Basics and Services Command Reference*

Chapter 2

# Unidirectional Links

This chapter describes unidirectional links and how to configure them. Topics include:



## Overview of Unidirectional Links

Most of the traffic in a broadcast video cable network is directed downstream to the user. Conventional bidirectional links do not optimize bandwidth allocation to match the bandwidth requirements of this mostly one-way traffic flow. In addition, the bidirectional nature of ports requires a port to receive data from the same port that it transmits data to. This behavior quickly consumes port resources without using them effectively.

You can conserve port resources and address the bandwidth requirements by implementing unidirectional links in the network.

Physical interfaces operate in bidirectional mode by default, both transmitting and receiving traffic. When you configure unidirectional mode on the interface, two new physical interfaces are automatically created. One interface, designated by -tx in the interface name, can only transmit traffic. The other interface, designated by -rx in the interface name, can only receive traffic. The parent physical interface is still present, but you effectively see a port with two unidirectional links. Figure 9 on page 27 illustrates the unidirectional nature of the new interfaces.

**Figure 9: Unidirectional Link Behavior**

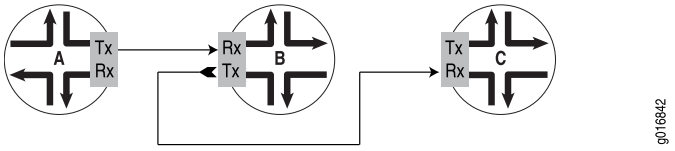


You can configure unidirectional mode on a per-port basis on the 10-Gigabit Ethernet interfaces of the 4-port 10-Gigabit Ethernet DPC on the MX960 Ethernet Services Router only. You can configure both unidirectional and bidirectional ports on a single DPC.

### *Configurable Options*

The transmit-only and receive-only interfaces created on a DPC port act independently. On the parent interface, you configure only the physical interface attributes common to both links. These attributes include clocking, framing, gigabit Ethernet options, and SONET options. On each of the unidirectional interfaces, you independently configure encapsulation (Ethernet only), MAC address, MTU size, address family (inet or inet6), and logical interfaces. VLAN tagging (untagged, single, stacked, or flexible) and VLAN IDs are also independently configurable on the receive-only and transmit-only interfaces. The full range of numbers for logical interfaces and VLAN IDs is available to both unidirectional interfaces.

To forward packets, you can configure only static ARP entries and static routes separately on each of the unidirectional interfaces. This configuration enables the transmit-only and receive-only interfaces to link to different ports on different routers. No other method of packet forwarding is currently supported.

The transmit-only and receive-only interfaces are removed when you delete unidirectional mode from the parent interface. The parent interface resumes operation as a normal, bidirectional interface.

### *Logical Interfaces*

You cannot configure logical interfaces on the parent interface after you have configured unidirectional mode. However, you can configure logical interfaces on both the transmit-only interface and the receive-only interface.

### *Alarm Reporting*

Alarms and defects are not reported for the transmit-only interface. Only local alarms and defects are reported for the receive-only interface. This behavior enables the use of SONET in a WAN-PHY configuration. SONET alarms, defects, and performance monitoring require bidirectional communication between sender and receiver. By accepting only local defects and alarms, the receiver interfaces in such a configuration are decoupled from the senders.

### *Operational State*

The transmit-only link on a unidirectional port is always operationally up. Operational state is not influenced by the state of the receive-only link on that port.

Operational state of the receive-only link on a unidirectional port is independent of the state of the transmit-only link on that port. Link state for a receive-only link is determined only by the status of locally detected faults on the that link. Change in the state of the receive-only link can trigger traps, flap messages, and alarms.

### *Statistics*

Statistics are reported differently for each of the three interfaces.

- Parent physical interface: No logical interfaces can be configured on the parent interface when it is in unidirectional mode. Therefore all traffic statistics for this

interface are reported as zero. All port-level statistics are reported on the parent physical interface rather than the rx or tx physical interfaces.

- Transmit-only physical interface: All transmit traffic statistics are reported for this interface. All receive (input) statistics are reported as zero. Also shown here are statistics for any logical interfaces configured on the transmit-only interface.
- Receive-only physical interface: All receive traffic statistics are reported for this interface. All transmit (output) statistics are reported as zero. Also shown here are statistics for any logical interfaces configured on the receive-only interface.

## System Requirements

To implement unidirectional links, you must use the following hardware and software components:

- JUNOS Release 8.5 or later for configuring unidirectional mode
- One or more MX960 routers
- One or more 4-port 10-Gigabit Ethernet DPCs installed in each MX960 router

## Configuring and Verifying Unidirectional Links

This section contains two examples and commands that you can use to configure and verify unidirectional links:


- Configuring and Verifying a Simple Example on page 29
- Configuring and Verifying a More Complex Example on page 31


### *Configuring and Verifying a Simple Example*

To configure unidirectional mode using default settings on 10-Gigabit Ethernet interface xe-5/1/0 and confirm the configuration:

```
[edit]
interfaces xe-5/1/0 {
    unidirectional;
}

[edit]
user@host show interfaces
xe-5/1/0 {
    unidirectional;
}
```

The transmit-only and receive-only interfaces are created as soon as you commit the configuration. The following show command is one way to verify creation of these new interfaces:

```
user@host run show interfaces xe-5/1/0* terse
Interface               Admin Link Proto    Local                 Remote
xe-5/1/0                up    down
```

```
xe-5/1/0-rx             up    down
xe-5/1/0-tx             up    up
```

The two unidirectional physical interfaces, xe-5/1/0-rx and xe-5/1/0-tx, are now present. In this example, no fiber-optic cables are connected to the port; consequently the xe-5/1/0 and xe-5/1/0-rx link states are down. In contrast, xe-5/1/0 is in the up state, because the transmit-only link is always up.

The following sample output provides more information about each of the interfaces. Unidirectional mode has been enabled on xe-5/1/0, the transmit-only and receive-only interfaces are present, and the link state matches expectations for no fiber-optic cables connected to the physical port.

```
user@host run show interfaces xe-5/1/0*
Physical interface: xe-5/1/0, Enabled, Physical link is Down
  Interface index: 318, SNMP ifIndex: 118
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Enabled, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:05:85:75:8b:62, Hardware address: 00:05:85:75:8b:62
  Last flapped   : 2007-08-10 11:45:29 PDT (01:39:47 ago)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0

Physical interface: xe-5/1/0-rx, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 129
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Rx-Only
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:05:85:75:8b:62, Hardware address: 00:05:85:75:8b:62
  Last flapped   : 2007-08-10 11:46:29 PDT (01:38:47 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0

Physical interface: xe-5/1/0-tx, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 130
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:05:85:75:8b:62, Hardware address: 00:05:85:75:8b:62
  Last flapped   : 2007-08-10 11:46:29 PDT (01:38:47 ago)
```

```
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
```

### *Configuring and Verifying a More Complex Example*

The following example makes the following changes from a default configuration:

- Sets framing mode to WAN-PHY on the parent interface, and consequently on the unidirectional interfaces as well.
- Configures VLAN IDs, VLAN tagging for single VLAN, and IP addresses on both unidirectional interfaces.
- Sets a nondefault MAC address on the receive-only interface.
- Configures a static ARP entry on the transmit-only interface. The entry contains a MAC address that is put into the Ethernet header destination address field of transmitted frames.

**CLI Quick Configuration**

To quickly configure the example described, copy the following commands and paste them into the router terminal window:

```
[edit]
set interfaces xe-5/1/0 framing wan-phy
set interfaces xe-5/1/0 unidirectional
set interfaces xe-5/1/0-rx mac 00:12:34:56:78:90
set interfaces xe-5/1/0-rx vlan-tagging unit 102 vlan-id 102 family inet address
  10.1.102.2/24
set interfaces xe-5/1/0-tx vlan-tagging unit 201 vlan-id 201 family inet address
  10.2.201.2/24 arp 10.2.201.3 mac 00:ab:cd:cd:ab:cd
set routing-options static route 10.33.1.1/32 next-hop 10.2.201.3
```

**Configuration Results**

Display the results of the configuration:

```
[edit]
user@host show interfaces
xe-5/1/0-rx {
    vlan-tagging;
    mac 00:12:34:56:78:90;
    unit 102 {
        vlan-id 102;
        family inet {
            address 10.10.102.2/24;
        }
    }
}
    xe-5/1/0-tx {
        vlan-tagging;
        unit 201 {
            vlan-id 201;
            family inet {
                address 10.2.201.2/24 {
                    arp 10.2.201.3 mac 00:ab:cd:cd:ab:cd;
                }
            }
        }
```

```
    }
    xe-5/1/0 {
      framing {
        wan-phy;
      }
      unidirectional;
    }
```

**Detailed Interface Information**

To display terse details about the interfaces:

```
user@host run show interfaces xe-5/1/0* terse
Interface               Admin Link Proto    Local                 Remote
xe-5/0/0                up    up
xe-5/0/0-rx             up    up
xe-5/0/0-rx.102         up    up   inet     1.1.102.2/24
                                  multiservice
xe-5/0/0-rx.32767       up    up   multiservice
xe-5/0/0-tx             up    up
xe-5/0/0-tx.201         up    up   inet     2.2.201.2/24
                                  multiservice
xe-5/0/0-tx.32767       up    up   multiservice
```

The additional logical interfaces for the unidirectional links result from the unit and VLAN tagging configuration.

To display more information about the interfaces:

```
user@host run show interfaces xe-5/1/0*
Physical interface: xe-5/1/0, Enabled, Physical link is Down
  Interface index: 151, SNMP ifIndex: 116
  Link-level type: Ethernet, MTU: 1514, Clocking: Internal, WAN-PHY mode, Speed:
 OC192, Unidirectional: Enabled, Loopback: None,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:05:85:75:8b:39, Hardware address: 00:05:85:75:8b:39
  Last flapped   : 2007-08-10 08:50:40 PDT (00:05:00 ago)
  Active alarms  : LOF, LINK
  Active defects : LOF, SEF, AIS-L, AIS-P, LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0

Physical interface: xe-5/1/0-rx, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 114
  Link-level type: Ethernet, MTU: 1518, WAN-PHY mode, Speed: OC192, Unidirectional:
 Rx-Only
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:12:34:56:78:90, Hardware address: 00:05:85:75:8b:39
  Last flapped   : 2007-08-10 08:50:40 PDT (00:05:00 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LOF, LINK
```

```
  Active defects : LOF, SEF, AIS-L, AIS-P, LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0

  Logical interface xe-5/1/0-rx.102 (Index 70) (SNMP ifIndex 115)
    Flags: Device-Down SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.102 ]  Encapsulation:
 ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.102/24, Local: 10.1.102.2, Broadcast: 10.1.102.255
    Protocol multiservice, MTU: Unlimited
      Flags: None

  Logical interface xe-5/1/0-rx.32767 (Index 71) (SNMP ifIndex 124)
    Flags: Device-Down SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation:
ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None

Physical interface: xe-5/1/0-tx, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 125
  Link-level type: Ethernet, MTU: 1518, WAN-PHY mode, Speed: OC192, Unidirectional:
 Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:05:85:75:8b:39, Hardware address: 00:05:85:75:8b:39
  Last flapped   : 2007-08-10 08:50:40 PDT (00:05:00 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface xe-5/1/0-tx.201 (Index 72) (SNMP ifIndex 126)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.201 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.2.201/24, Local: 10.2.201.2, Broadcast: 10.2.201.255
    Protocol multiservice, MTU: Unlimited
      Flags: None

  Logical interface xe-5/1/0-tx.32767 (Index 73) (SNMP ifIndex 127)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

You can use the show interfaces xe-5/1/0* extensive command to display the most complete set of information about the interfaces. Alternatively, you can specify only xe-5/1/0, xe-5/1/0-rx, or xe-5/1/0-tx to show extensive information about just one interface.

The extensive output includes statistics for the interfaces. The following excerpts show the differences between the receive-only and transmit-only interfaces for statistics.

In the following output for a receive-only interface, input statistics are recorded, but all output statistics have a value of zero.

```
user@host show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :      322857456303482           9627496104 bps
   Output bytes  :                   0                    0 bps
   Input  packets:         328775413751              1225495 pps
   Output packets:                   0                    0 pps

...

  Filter statistics:
    Input packet count            328775015056
    Input packet rejects                     1
    Input DA rejects                         0

...

  Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :      322857456303482
     Output bytes  :                   0
     Input  packets:         328775413751
     Output packets:                   0

...

Transit statistics:
     Input  bytes  :      322857456303482           9627496104 bps
     Output bytes  :                   0                    0 bps
     Input  packets:         328775413751              1225495 pps
     Output packets:                   0                    0 pps

...
```

In the following output for a transmit-only interface, output statistics are recorded, but all input statistics have a value of zero.

```
user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
```

```
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                   0 bps
   Output bytes  :      322891152287160          9627472888 bps
   Input  packets:                    0                   0 pps
   Output packets:         328809727380             1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0

...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :      322891152287160
     Input  packets:                    0
     Output packets:         328809727380

...

Transit statistics:
     Input  bytes  :                    0                   0 bps
     Output bytes  :      322891152287160          9627472888 bps
     Input  packets:                    0                   0 pps
     Output packets:         328809727380             1225492 pps

...
```

**Related Topics**

For more information about concepts associated with unidirectional links, see the following resource:

- RFC 3077, *A Link-Layer Tunneling Mechanism for Unidirectional Links*

# Part 2
# Voice Solution

Chapter 3

# Overview of the Voice Solution

This chapter describes the Juniper Networks voice solution. Topics include:



## The Voice Solution in a Next-Generation Network Overview

The voice solution provides a way for the router to integrate into a Telecommunications and Internet Converged Services and Protocols for Advanced Networks (TISPAN)/IP multimedia subsystems (IMS) environment to provide voice over IP (VoIP) functionality. IMS is a flexible network architecture that allows providers to introduce multimedia services across both next-generation packet-switched and traditional circuit-switched networks. It uses open interfaces and functional components that can be assembled flexibly to support real-time interactive services and applications.

IMS provides a standards-based architecture that allows mobile carriers to migrate to next-generation networks that support applications that combine voice, video, and data functionality. The European Telecommunications Standards Institute (ETSI) created TISPAN to extend IMS support to fixed-line carriers. This extension is commonly called fixed mobile convergence (FMC). IMS/FMC allows subscribers to

access any network (wireless or fixed) from any device (computer, PDA, or cell phone) and to move seamlessly from one network to another.

The router acting as a packet gateway provides much of the border gateway function (BGF), as shown in the ETSI-TISPAN architecture in Figure 10 on page 40:

**Figure 10: Routers Running JUNOS Software in the ETSI-TISPAN Architecture**

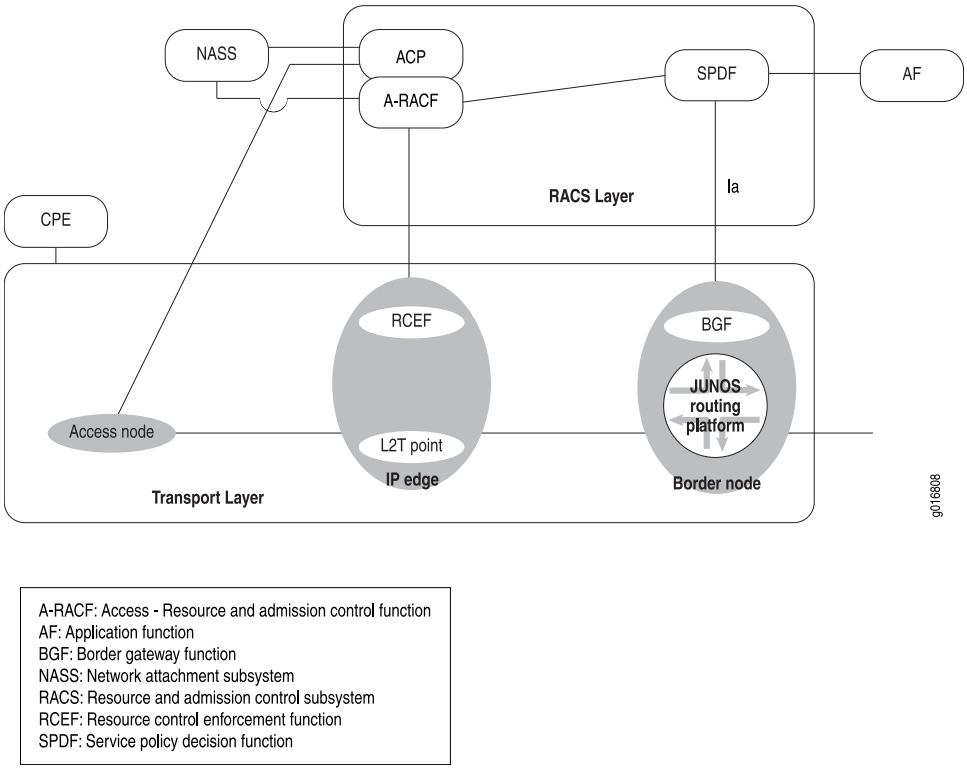


A-RACF: Access - Resource and admission control function
AF: Application function
BGF: Border gateway function
NASS: Network attachment subsystem
RACS: Resource and admission control subsystem
RCEF: Resource control enforcement function
SPDF: Service policy decision function

## *Voice Solution Terms and Abbreviations*

Table 8 on page 40 defines the terms and abbreviations used in this topic.

**Table 8: Terms and Abbreviations**

| Term | Description |
|---|---|
| BGF | Border gateway function. Resides in the transport layer and polices and enforces traffic flows based on instructions from the SPDF. The packet gateway provides the BGF functionality. |
| Context | An association between terminations. |
| Gate | Unidirectional flow of IP packets as directed by the PGC. Sometimes called a pinhole. |
| Ia | A profile of the interface between an SPDF (the PGC) and the BGF (packet gateway). |
| I-BGF | Interconnect-BGF. The BGF between two peering partners. The packet gateway in the router provides much of the I-BGF function. |

**Table 8: Terms and Abbreviations** *(continued)*

| Term | Description |
|---|---|
| IMS | IP multimedia subsystem. |
| IPC | The VPGs and the MultiServices PIC communicate by exchanging Inter-Process Communication (IPC) messages over a TCP connection; this is internal (intra-chassis) communication. |
| PG | Packet gateway. A virtual device on the router that provides media processing and control as directed by the PGC. |
| PGC | Packet gateway controller. An external device that provides signal processing and directs the behavior of the PG. The PGC provides the service policy decision function (SPDF) shown in Figure 10 on page 40. |
| PGCP | Packet Gateway Control Protocol (PGCP). An H.248 v3 protocol with Juniper Networks extensions. It provides management and signaling between the PG and the PGC. |
| pgcpd | The packet gateway pgcpd process running in the Router Engine decodes H.248 messages that VPGs receive from the PGC and translates the H.248 messages to IPC messages. |
| SPDF | Service policy decision function. Controls the BGF. In the Juniper Networks voice solution, the PGC acts as the SPDF. |
| Stream | A bidirectional flow within a context. |
| Termination | A local source and sink of packets. |
| Virtual packet gateways (VPGs) | A VPG consists of a packet gateway configuration on the Routing Engine. VPGs are controlled by a PGC. A VPG receives instructions from the PGC and instructs the MultiServices PIC how to treat voice traffic. |

## Voice Solution Architecture

As shown in Figure 11 on page 42, the two main components of the voice solution are the packet gateway controller (PGC) and the packet gateway (PG). The PGC and the PG communicate over the Packet Gateway Control Protocol (PGCP).

**Figure 11: Voice Solution Architecture**

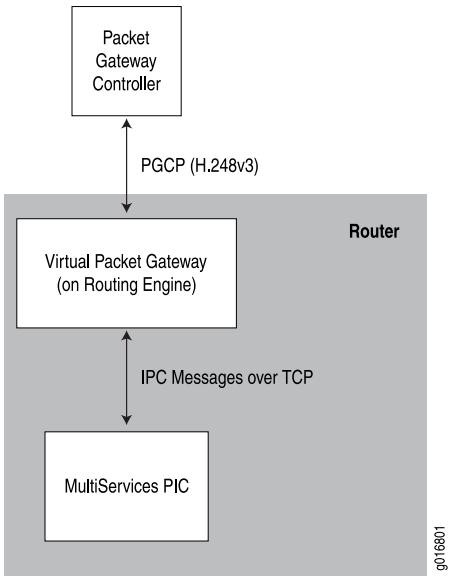


## *Packet Gateway Controller*

The PGC is an external device that controls the PG on the router. The PGC requests media services and resource allocation from the PG, and it uses those services and resources for VoIP call signaling setup. The PGC maintains awareness and control over the network's transport resource using PGCP connections with all of the PGs in the network.

## *Packet Gateway*

The packet gateway feature on the router provides Interconnect-BGF transport services for VoIP sessions. The packet gateway feature consists of:

- VPGs
- pgcpd process
- MultiServices PIC that controls voice traffic based on instructions it receives from the VPG.

## *PGCP*

The PG and the PGC communicate over a Packet Gateway Control Protocol (PGCP) connection. PGCP is an H.248 v3 protocol with Juniper Networks extensions. PGCP is compliant with *Gateway control protocol v3, ITU T Recommendation H.248.1, September 2005*.

## Voice Solution Topology with Multiple VPGs and PGCs Overview

The voice solution supports up to eight concurrent VPGs in a router. Each VPG is connected to a PGC over its own PGCP connection. One VPG can connect to one PGC at the same time. Multiple VPGs can share a single service PIC. A single VPG cannot span more than one service PIC.

Creating multiple VPGs allows you to deploy different policy and quality of service (QoS) characteristics in your network. It also allows you to scale your infrastructure by using multiple MultiServices PICs to control voice traffic.

Figure 12 on page 43 shows a topology with multiple VPGs and PGCs. This topology allows one VPG and one MultiServices PIC to continue handling gate requests and forwarding packets on open gates even when the other PIC fails.

**Figure 12: Topology with Multiple VPGs and PGCs**

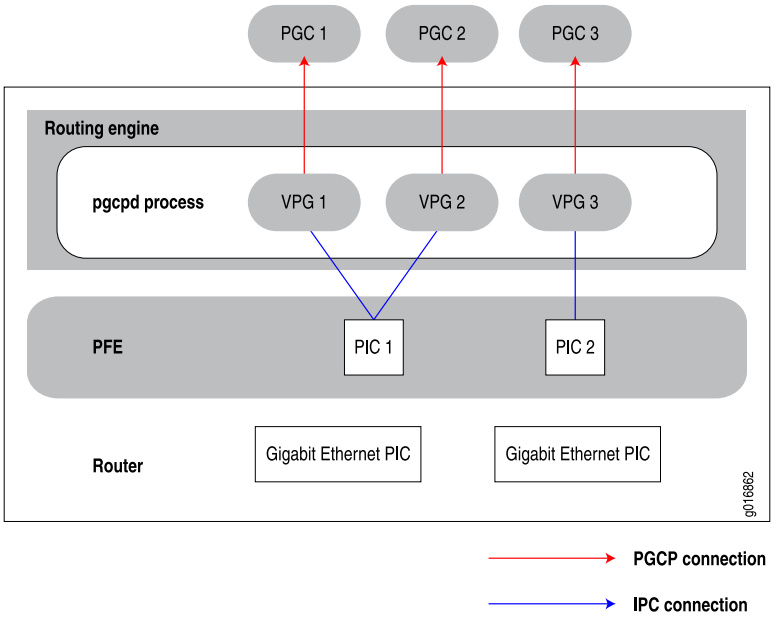


You can have multiple PGCs configured for one VPG. When a VPG begins running on the router, it attempts to set up a connection to the first configured PGC. Each VPG can have one active PGC and one or more standby PGCs. In case of a PGC failure or in case of the PGC sending instructions to the VPG, the VPG can switch to another PGC. Figure 13 on page 44 shows an active and standby PGC connected to VPG 2.

**Figure 13: Active and Standby PGCs**

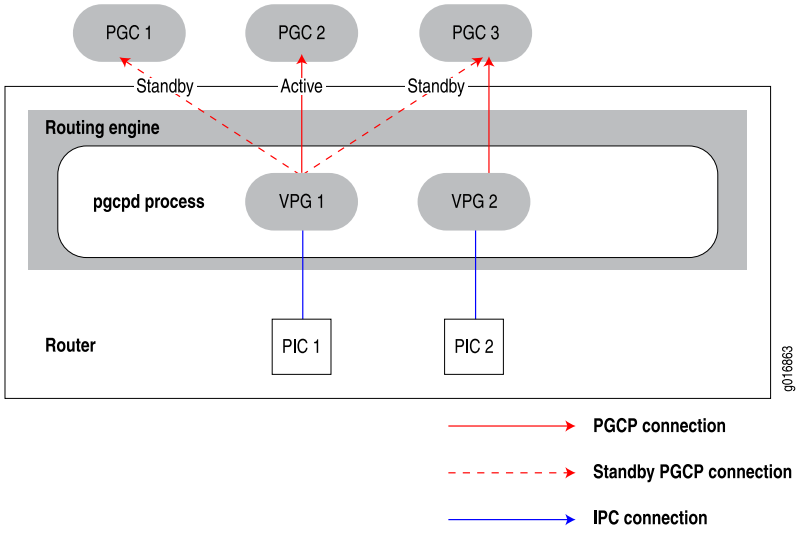


If the PGCP connection between the VPG and the PGC is lost, the VPG attempts to reconnect to the PGC. If the VPG cannot reconnect to the PGC, it traverses its list of PGCs until it successfully connects to one of the PGCs.

## Sample Voice Network Topology

Figure 14 on page 44 shows a sample network that uses the voice solution.

**Figure 14: Sample Voice Network**

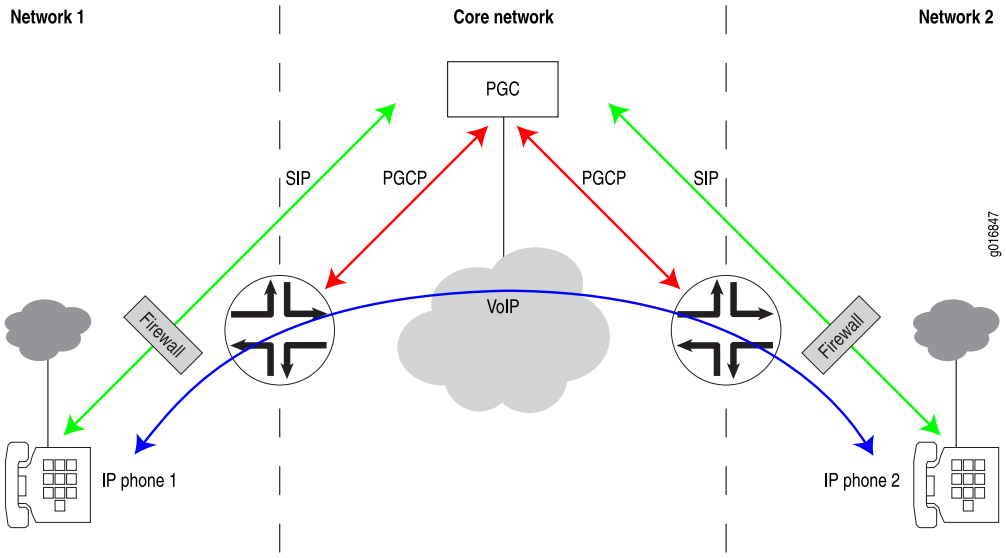

## Control of Voice Flows with Gates Overview

The voice feature uses gates to control voice flows in the transport plane. Gates are created through signaling instructions that the PGC provides to the PG. Using the signaling instructions, the PG defines gates to allow, drop, or manipulate voice flows as they traverse the router.

Each gate provides a unidirectional voice flow. A pair of gates provides a bidirectional voice flow. Figure 15 on page 45 shows a unidirectional gate.

**Figure 15: Unidirectional Gate**

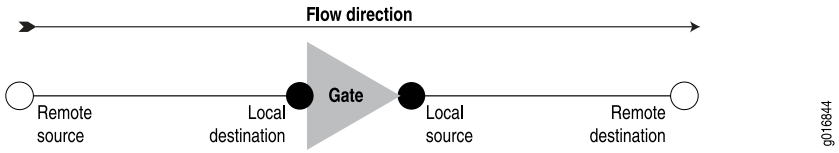


### Gate Addressing

Gates are defined by their local source and destination addresses and their remote source and destination addresses.

Figure 16 on page 45 shows a gate pair, which represents a bidirectional voice flow. The local destination address of Gate 1 is equal to the local source address of Gate 2, and the local source address of Gate 1 is equal to the local destination address of Gate 2.

**Figure 16: Addressing of Gate Pairs**

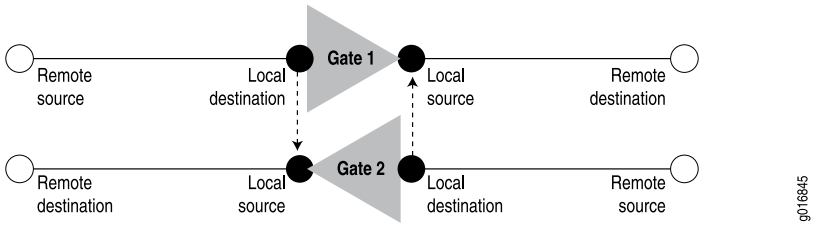


### Opening, Closing, and Modifying Gates Overview

Based on information acquired through VoIP signaling, the PGC instructs the packet gateway through PGCP commands which gates to create and which actions to associate with them. Each gate can have many actions associated with it; for example, NAT, Differentiated Services (DiffServ) code point (DSCP) marking, and latching. The pgcpd process decodes PGCP commands that it receives from the PGC and uses IPC messages to instruct the PIC to create, delete, or modify gates and apply required actions to each gate.

The following IPC messages are exchanged between the pgcpd process and the PIC:

- Gate open request
- Gate close request
- Gate audit request
- Gate modify request
- Gate open reply
- Gate close reply
- Gate audit reply
- Gate modify reply
- Gate notification reply

### *Gate Identification*

When a gate is created, it is assigned an identifier. You can use this identifier with the PGCP show commands to monitor specific gates.

### *Forward and Drop Operations for RTP and RTCP Gates*

You can use the StreamMode property in the LocalControl Descriptor of H.248 messages to change the mode of Realtime Transport Protocol (RTP) gates without affecting the mode of Real-Time Control Protocol (RTCP) gates. That is, you can put RTP gates in drop mode while leaving RTCP gates in forward mode.

To view whether RTP and RTCP gates are in drop mode or forward mode, use the show services pgcp flows command. The following example shows a gate in which the RTP stream is in drop mode, and the RTCP stream is in forward mode.

```
user@host>show services pgcp flows
Gate id: 4295033089
UDP      20.50.170.110:0     ->    20.50.170.2:1024 Drop I              0
    NAT source   20.50.170.110:0       ->     10.50.170.1:1024
    NAT dest       20.50.170.2:1024    ->   10.50.170.110:20000
Gate id: 4295033089
UDP      20.50.170.110:0     ->    20.50.170.2:1025 Forward I           0
    NAT source   20.50.170.110:0       ->     10.50.170.1:1025
    NAT dest       20.50.170.2:1025    ->   10.50.170.110:20001
```

### *Latch Deadlock and Media Inactivity Detection and Reporting*

You can configure the parameters that a VPG uses to detect and report latch deadlocks.

**Detection**—The VPG uses an inactivity timer to detect a latch deadlock or other media inactivity on a gate. The timer tracks the receipt of media packets during a specified time interval.

When a latching signal exists for a termination, the PG places the termination in a Drop state. All incoming traffic to the relevant gate egressing the termination is dropped until the first IP traffic datagram enters the termination (ingress). At this point the remote descriptor on the termination and egress gate is updated to forward

traffic to the newly acquired source. The latch signal is removed from the gate when the PGC receives an H.248 Notify message containing the newly acquired IP address. Deadlock occurs when an error occurs regarding the source IP address and port that prevents the endpoint from returning data to the source address.

The detection process is activated in one of the following ways:

- The PGC requests quality (QUA) alerts.
- The PGC requests application data inactivity detection (ADID) alerts.
- The CLI is used to configure a forced service change when media inactivity occurs.

The inactivity timer tracks the receipt of media packets during a specified interval of time (inactivity duration). If no media packets are received during this time interval, the VPG reports the inactivity to the PGC.

**Reporting**—Reporting of the media inactivity occurs in one of the following ways:

- If the VPG detects a latch deadlock or media inactivity and you have configured the VPG to force a service change, the VPG stops service on the gate and sends the PGC a ServiceChange message using either error code 906 (Loss of Lower Layer Connectivity) or 910 (Media Capability Failure). The VPG then takes the affected termination out of service, but does not subtract the termination.
- If you have not configured the VPG to force a service change, the VPG sends a QUA or ADID notification message, depending on which type of notification was requested by the PGC.

**Related Topics**

- Configuring Latch Deadlock and Media Inactivity Detection on page 74
- Monitoring Gates on page 87

## H.248 Building Blocks Overview

The H.248 connection model uses contexts, terminations, and streams, which are logical entities that the PGC controls. In the router, the MultiServices PIC creates a context. The software then adds terminations to the context and adds streams to the terminations. Figure 17 on page 48 shows a context, termination, and stream.

**Figure 17: Context, Termination, and Stream**

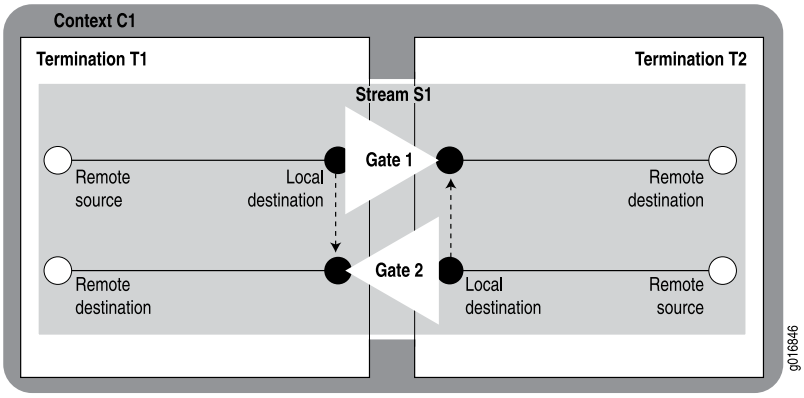


### *Terminations*

A termination can be a source and sink for media and control streams, and the parameters of the streams are encapsulated within the termination. A termination is characterized by properties that are grouped in a set of descriptors that are included in add, subtract, modify, or audit commands. Terminations have unique identifiers (TerminationIDs) that the packet gateway assigns when it creates the termination.

Each termination is the source and destination of a gate. A termination exists only as long as a call. It is removed when the call is removed.

### *Contexts*

A context is an association between a collection of terminations. The VPG instructs a MultiServices PIC to create a context for each voice session and each signaling session. Using instructions from the VPG, the PIC then applies policies such as DSCP, NAT, rate limiting, and inactivity timers to the gates within a context. If the VPG does not specify an existing context to which the termination is to be added, the PIC creates a new context.

### *Streams*

A stream is one bidirectional flow within a context.

## Virtual Interfaces with the Packet Gateway Overview

The packet gateway and the PGC communicate through virtual interfaces. JUNOS interface names are not known or communicated to the PGC. You configure a virtual interface on the packet gateway, and this virtual interface is provided to the PGC. The virtual interface configuration includes the media service for the virtual interface, which contains the name of the NAT pool.

Included in the H.248 message exchange between the PGC and the VPG is a virtual interface identifier. This identifier instructs the packet gateway which media resources to use.

**Related Topics**

- Configuring Virtual Interfaces on page 65
- Example: Configuring the Virtual Interfaces on page 151

## Twice NAT for VoIP Traffic Overview

The packet gateway supports both network address translation (NAT) and network address port translation (NAPT). *Twice NAT* enables you to configure both source addresses and destination addresses that are translated as packets traverse the router. You can apply twice NAT for VoIP packets (signaling and media) as they traverse gates to achieve security between realms or service providers. To apply twice NAT, the pgcpd process instructs the PIC to allocate a specified number of NAT addresses and ports from a PGCP NAT pool on a per-gate basis. The pgcpd process specifies which NAT pool to use.

Figure 18 on page 49 shows two gates in a packet gateway.

**Figure 18: Translation of Gate Addressing**

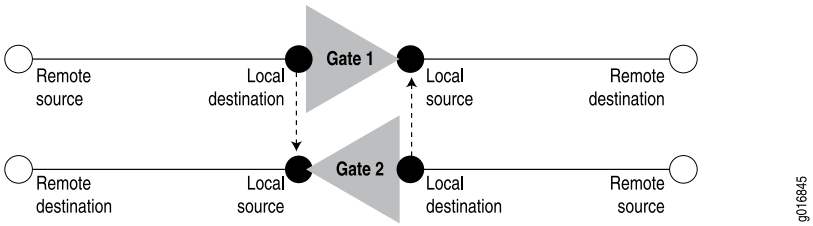


After flows are created for Gate 1, the gate connects the remote source to the local destination. The local source and local destination addresses reside on the router and must be uniquely specified. For Gate 1, twice NAT enables the router to translate the IP address of the remote source to the local source, and the local destination to the remote destination.

To create the bidirectional flow, the same IP address is used for the local source in Gate 2 and the local destination in Gate 1. Likewise, the same IP address is used for the remote source in Gate 1 and the remote destination in Gate 2.

Figure 19 on page 49 shows an example of how addresses are translated.

**Figure 19: Example: Translation of Gate Addressing**

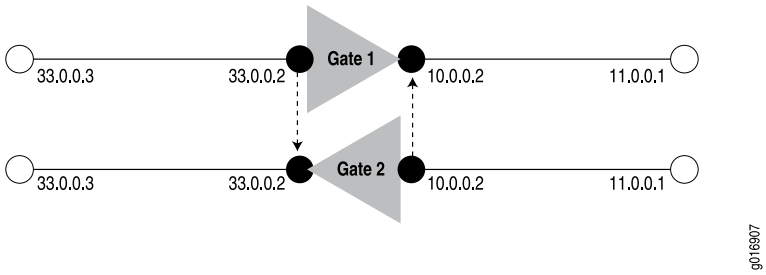

## *NAT Pool Selection*

You can configure separate NAT pools that can be controlled by either the PG or the PGC. By default the PG controls the addresses and ports in a pool. However, when you configure your NAT pool, you can specify that the PGC controls the addresses and ports in the NAT pool. The PGC reserves the addresses and ports when it requests specific local NAT bindings for remote addresses.

If the PG selects the NAT pool, it can use one of the following methods to select the pool:

- (Default) Using the value of the media services assigned to virtual interfaces configured on the PG.
- Matching the transport protocol type in H.248 messages received from the PGC.

### NAT Pool Selection by Matching the Transport Protocol

The PG can select the NAT pool by matching any combination of the following protocols:

- Real-Time Transport Protocol using Audio/Video profile (RTP/AVP)
- TCP
- UDP

Selecting a NAT pool based on transport protocol:

- Guarantees the prioritized distribution of network resources.
- Enables the use of multiple NAT pools for each virtual interface.

The PGC can set a transport protocol in the media description in the local descriptor command in Add and Modify commands that it sends to the PG. The media description format is:

```
m=<media> <port> <transport> <format list>
```

where the transport field specifies the transport protocol. For example:

```
m=video 49170/2 RTP/AVP 31
```

When you set up your NAT pools, you specify a transport protocol or list of protocols. Do not configure the NAT pool to be remotely controlled by the PGC. Also, set the port in the NAT pool to automatic.

When the PG receives an Add or Modify command with a media description, it searches the NAT pools associated with the virtual interface and attempts to match the transport protocols in the description with the transport protocols specified in the NAT pools. The PG uses the first NAT pool that has a matching transport protocol. If it cannot find a match, it replies to the PGC with the following error:

```
ER=500 {"Application: Media handler not found"}
```

### *IPv4-to-IPv6 Address Translation*

IPv4-to-IPv6 address translation enables callers in an IPv4 network to place calls to recipients in an IPv6 network. With this capability, the access side of the network can be an IPv4 network and the backbone side of the network can be an IPv6 network and vice versa. The PGC sets up gates so that one termination of the gate has IPv4 addresses and the other termination of the gate has IPv6 addresses. The packet gateway performs the appropriate IPv4-to-IPv6 and IPv6-to-IPv4 translations.

This implementation is not the tunnelling of IPv4 headers over IPv6 headers and vice versa. It is the translation of the IPv4 headers to IPv6 headers and vice versa.

You must configure both an IPv4 NAT pool and an IPv6 NAT pool on the PG for IPv4-to-IPv6 translation to work.

Figure 20 on page 51 shows an example of a gate pair in a network where IPv4-to-IPv6 address translation is used.

**Figure 20: IPv4-to-IPv6 Gates Using Twice NAT**

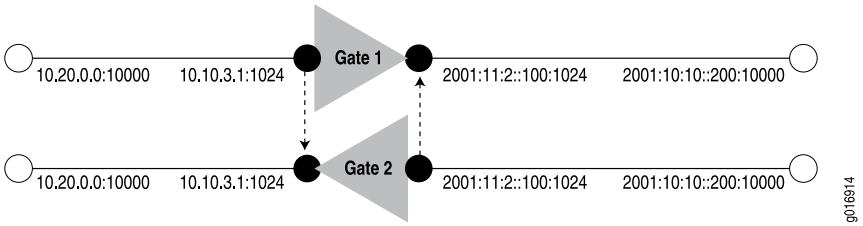


**Related Topics**

- Configuring NAT Pools for the Packet Gateway on page 63
- Assigning a NAT Pool on page 65
- Example: Configuring NAT Pools for the Packet Gateway on page 146
- Example: Assigning the NAT Pools to a Media Service on page 150

## Quality of Service for VoIP Traffic Overview

To ensure optimized quality conditions for VoIP traffic, in gate open requests, the PGC can include a request for the packet gateway to mark voice traffic with various DSCP code points. The pgcpd process passes this information to the MultiServices PICs, which then apply these actions to the gate.

You can configure a default DSCP value that the VPG uses for outgoing traffic when the DSCP value is not defined by the PGC. If you do not configure a value, the default value is 0x00. All eight bits are exposed, but the packet uses only the six leading bits. You can embed other data in the other two bits.

The DiffServ package is defined in Annex A.2 of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

**Related Topics**

- Configuring QoS for Voice Calls on page 70

■ Example: Configuring QoS for Voice Calls on page 157

## Rate-Limiting for VoIP Traffic Overview

Because PGCP traffic flows involve voice traffic, the flows require quality of service that:

- Provides the bandwidth that the flow requires.
- Ensures that flows do not consume more resources than they need.
- Regulates flows that are nonconforming and present vastly greater rates of traffic.

The packet gateway provides a two-rate policer that you can apply to the ingress traffic of any gate.

This quality of service is provided through a two-rate three-color policing functionality on the MultiServices PIC. This policer complies with *RFC 2698, A Two Rate Three Color Marker, September, 1999*. With the rate limiting capability, the MultiServices PIC can police flows to conform to:

- Committed Information Rate (CIR)
- Peak Information Rate (PIR)
- Committed Burst Size (CBS)
- Peak Burst Size (PBS)

### *How the Rate-Limiting Feature Works*

You use rate limiting with gates. To enable rate limiting for a gate, you need to provide traffic management package (TMAN) parameters. You can configure these parameters in the CLI or they can come from the PGCP signaling commands received from the PGC. Traffic-management parameters that come from the PGC override parameters configured in the CLI.

**Table 9: Traffic Parameters Configured in the CLI**

| Parameter | Description | Equivalent PGCP Signaling Command |
|---|---|---|
| Sustained data rate (SDR) | Provides the CIR | Tman/sdr |
| Peak data rate (PDR) | Provides the PIR | Tman/pdr |
| Maximum burst size (MBS) | Provides the burst size. Both the CBS and the PBS defined in RFC 2698 map to the maximum burst size. | Tman/mbs |

For each of the traffic-management parameters, you can configure a default value that applies to all gate streams and value that applies only to RTCP gate streams. For RTCP streams, you can specify a fixed value for the parameters or you can specify

the value as a percentage of the RTP rate. When RTP and RTCP are represented as a single stream, RTCP is policed whenever RTP is policed.

The PGC can send traffic-management parameters to the packet gateway in PGCP gate open and gate modify signaling requests. When the services PIC receives these parameters, it marks the packets red, yellow, or green as specified in RFC 2698. A packet is marked red if it exceeds the PIR. A packet is marked yellow if it exceeds the CIR. A packet is marked green if it does not exceed the CIR. Packets that are marked red are dropped by the services PIC.

#### Default Values for Rate-Limiting Parameters

If the policy command H.248 message from the PGC is on (tman/pol = on), but the rate-limiting parameters are not specified in the message and the JUNOS rate-limiting parameters have not been configured, the PG uses following default values:

- Peak data rate—10,000 bytes per second for all streams and 5 percent of the RTP gates' PDR for RTCP streams.
- Sustained data rate—10,000 bytes per second for all streams and 5 percent of the RTP gate's SDR for RTCP streams.
- Maximum burst size—1000 bytes for all streams and the MBS of the RTP gate for RTCP streams.

### *Rate Limiting and Fast Update Filters*

When a VoIP flow configured through the packet gateway violates the SDR by three times the configured rate, fast update filters are installed on the gate to allow the rate-limiting drop action to occur on the PFE instead of the PIC.

A fast update filter is similar to a regular filter that is defined in the [edit firewall] hierarchy, except that the system can incrementally add or update terms.

For fast update filters, a term equals a gate definition. You can see gate definitions in the show services pgcp extensive command output.

The fast update filter match is performed based on the most specific defined term. For each filter, a default term is installed to allow traffic to pass through (otherwise, all traffic is dropped because it is the default firewall action). For example, two terms are listed when there are two filters.

Filters are in effect until the gate is destroyed. If the client loses its connection for over 30 seconds, the existing filters are deleted, and default fast update filters are installed.

### *Rate-Limiting Statistics Display*

To display statistics for a gate including rate-limiting statistics and the number of packets dropped because of FUF filters, use the show services pgcp gates gate-id *gate-id* statistics command.

**Related Topics**

- Configuring Rate-Limiting for Voice Calls on page 69

■ Collecting Statistics on Gates with Rate-Limited Flows on page 90

## Security for PGCP Overview

The PGCP feature provides the following security features:

- Interim AH scheme
- Symmetric control association

### *Interim AH Scheme*

If the underlying network layer does not support IPSec, you can use the interim authentication header (AH) scheme to provide security on the connection between the VPG and the PGC. The interim AH scheme defines an authentication header with the H.248 protocol header.

To use the interim AH scheme, configure the security algorithm for the interim AH scheme for a PGC. If you configure an algorithm, the PG accepts H.248 messages from the PGC that include an AH from the defined algorithm. It discards received packets that do not include the expected AH. When the PG replies to the PGC, it includes an AH from the defined algorithm.

### *Symmetric Control Association*

For control association between the PG and a PGC, you define the address and port of the PG and the PGC. The PG uses the address and port configured for the PGC when it sends registration messages to the PGC. If the registration reply contains a ServiceChangeAddress command, the PG connects to the PGC using the new address or port or both instead of the address and port configured in the CLI. The PG accepts only H.248 messages that arrive from the PGC address and port. All other messages are dropped.

In the following cases, the PG attempts to connect to the address and port configured on the router:

- Loss of the PG-to-PGC connection
- Restart of the pgcp-services
- Reboot of the router

If needed, the PGC can reply with a new ServiceChangeAddress command.

The PG uses the new address in the ServiceAddressChange command only if the command is triggered by ServiceChangeReason 901 & 902. If the change is triggered by other ServiceChangeReasons such as 900, the PG uses the configured address and port.

**Related Topics** ■ Adding a PGC to the VPG Configuration on page 62

## Priority and Emergency Call Handling

The PGC can set values for priority and emergency indicators for a context and include them in add or modify requests that it sends to the packet gateway. The packet gateway stores these priority and emergency values. The PGC can then query the packet gateway for context lists based on the priority and emergency settings. The packet gateway includes the priority and emergency properties in add or modify commands when the PGC requests a ContextAudit.

The packet gateway does not provide higher queuing and processing for emergency and priority calls.

## VoIP Call Setup Overview

As shown in Figure 21 on page 55, VoIP uses two streams:

- Signaling stream, which handles the agreement to set up calls. The signaling stream can use Session Initiation Protocol (SIP) or other signaling protocols.
- Media (RTP/RTCP) stream for each leg of the voice call.

**Figure 21: Establishing a VoiP Call**

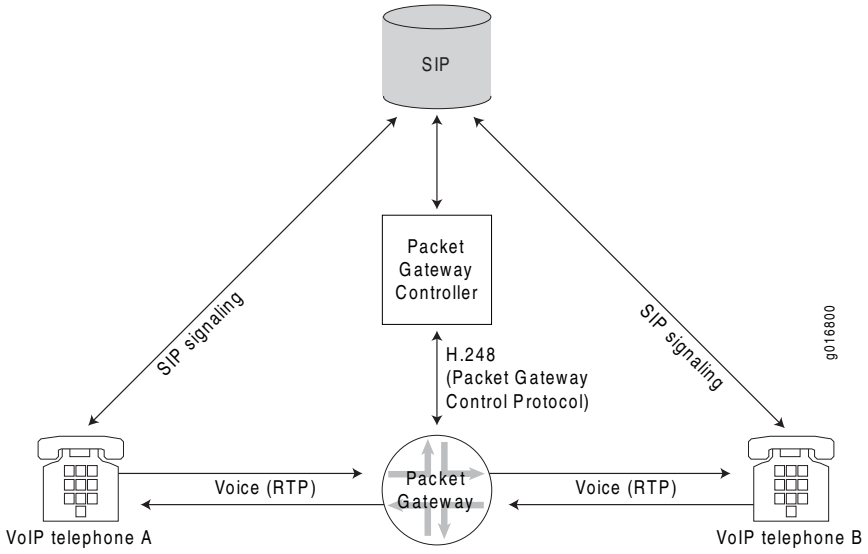


The process of setting up a VoIP call in the network using SIP, as shown in Figure 21 on page 55, is as follows:

1. VoIP telephone A initiates a VoIP call to VoIP telephone B.
2. VoIP telephone A sends a SIP message to the SIP server.
3. The PGC (SIP server) sends an H.248 request for gate allocation from the VPG.

4. The pgcpd process running on the Routing Engine sends IPC messages to the PIC requesting that the PIC open gates for each call leg.
5. The PIC creates the gates with the behaviors specified in the IPC messages, and it sends a reply to the pgcpd process. Gates are allocated in a Drop state.
6. The VPG sends an H.248 response providing allocated gate information to the PGC.
7. The SIP server sends the modified SIP signaling (based on the gate info sent by the VPG) to the destination VoIP telephone B.
8. VoIP telephone B replies to the SIP request to the SIP server.
9. The PGC updates the VPG with the new information sent by VoIP telephone B.
10. Steps 4-6 are repeated, where the PIC is updated with the new information provided by the PGC. Gates are transitioned into a Forward state
11. The SIP server sends the modified reply to VoIP telephone A.
12. The call is established. Media streams can now flow through the routers' open gates.

## VPN Aggregation for VoIP Calls Overview

The VPN aggregation feature uses VPN routing and forwarding (VRF) so users on one VPN can call users on another VPN. For example, in Figure 22 on page 57, users in VPN B can call users in VPN A and VPN C.

**Figure 22: VPN Aggregation in a VoIP Network**

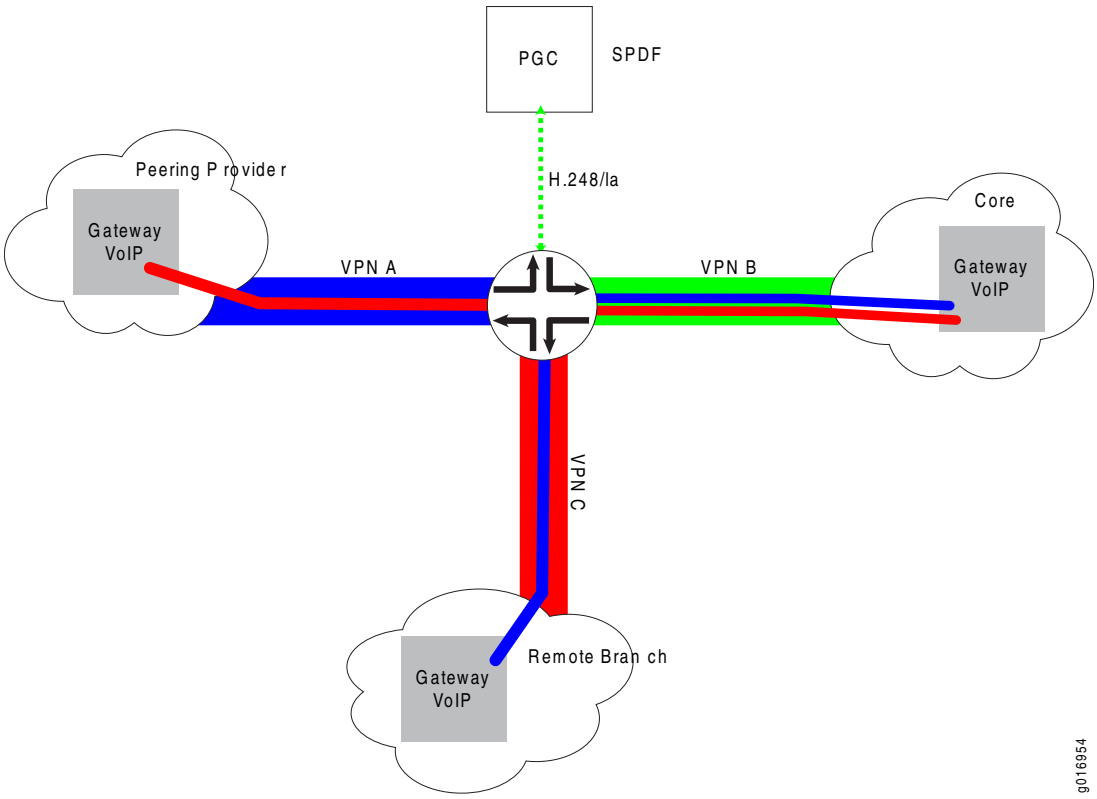


VPN aggregation provides the following benefits:

- Provides a scalable way to configure VRFs in a mesh-like configuration that uses only one logical service interface for each VRF.
- Reduces the number of service sets that you need because you can add all of your logical service interfaces to a pool of interfaces, and then assign the entire pool of interfaces to a service set.
- Configurations are inline so, when you provision the service set for VRFs, you can seamlessly tie the service into the PGCP service without the need for additional configuration states.
- Uses the router's native support for VRFs and VPNs, which omits the need for an external element that terminates the VRFs and replaces them with the VLAN tags required to support VoIP media handling.

## *How VPN Aggregation Works*

VPN aggregation uses the virtual interface configurations as shown in Figure 23 on page 58 to route traffic from users in one VPN to users in another VPN.

**Figure 23: Overview of VPN Aggregation Configuration**

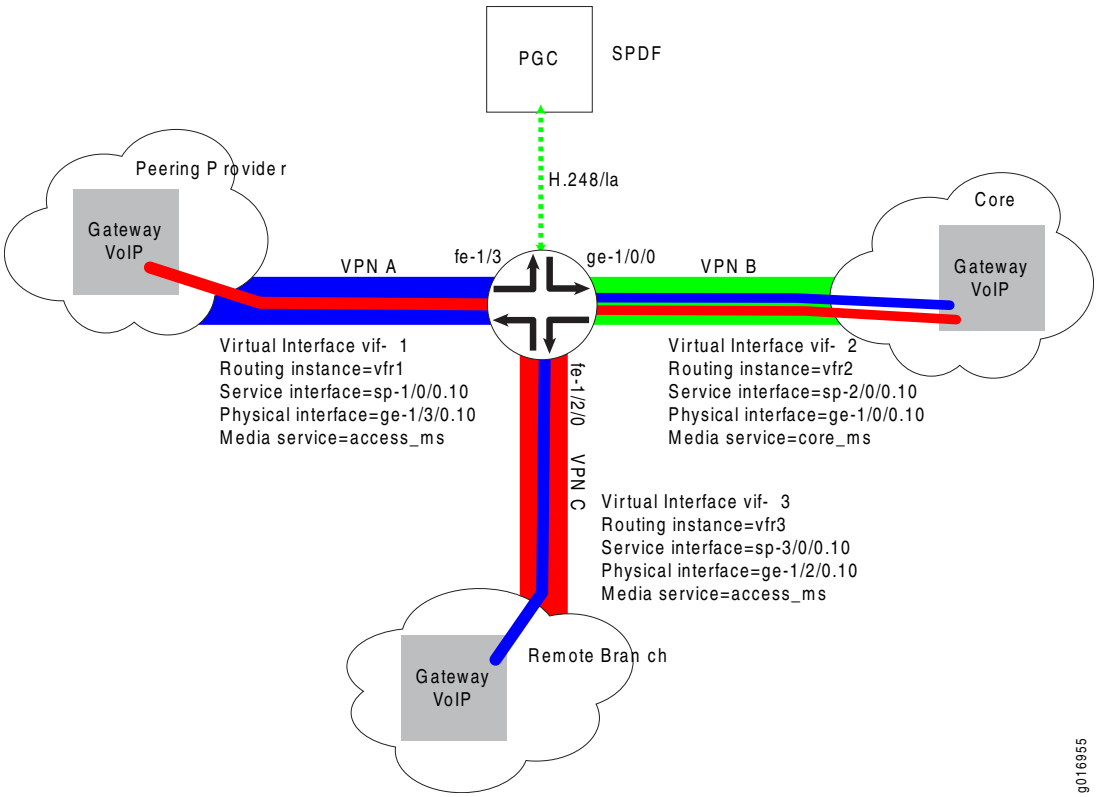


The VPN aggregation configuration consists of:

- VRFs—One for each VPN. The VRF is required to create a layer 3 VPN. The VRF must have the instance type of VRF, a physical interface and logical service interface, a route distinguisher, and VRF import and export policies.
- Pool of logical service interfaces—One pool that contains all service interfaces that are configured in your VRF routing instances. Instead of explicit inside and outside service interfaces, all of the interfaces in the pool can be both inside and outside service interfaces.
- Service Set—One service set that has a next-hop service set to the pool of logical service interfaces and that contains a PGCP rule. The service set links the VRFs to the PGCP service.
- Virtual interface—One for each VRF routing instance. The virtual interface configuration establishes the relationship between the following parts of the configuration:
  - PGCP NAT pool (the media service contains the NAT pool)
  - VRF routing instance to which the NAT routes are added
  - The service interface
  - The physical interface

When a gate is established, the pgcpd process uses the virtual interface information in the termination ID to determine the ingress and egress virtual interfaces for the gate. In turn, the virtual interface configuration maps to the VRF, NAT pool, service interface, and physical interface.

The termination IDs of the caller and the call recipient contain the virtual interface ID. For example, in Figure 23 on page 58 termination ID ip/4/vif-1/1 matches virtual interface vif-1, which is mapped through the configuration to routing instance vrf1.

**Related Topics**

- Configuring VPN Aggregation on page 72

## Session Mirroring Overview

Session mirroring allows you to send a copy of a context to an external device called a delivery function for analysis. With session mirroring, the original session is sent to its intended destination and the mirrored session is sent to the delivery function. The mirroring operations are transparent to the user whose session is being mirrored.

Session mirroring is supported only for IPv4 traffic.

The packet gateway can mirror up to 1 percent of gates at a time.

### *Activation of Session Mirroring for a Gate*

When session mirroring is enabled, the PG uses information in PGCP requests received from the PGC to identify sessions to be mirrored and to trigger the mirroring session. The following sample PGCP request includes session-mirroring information:

```
MEGACO/2 [123.123.123.3]:2944
Transaction = 10003 {
    Context = $ {
        Add = $ {
            Media {
                LocalControl {
                    Mode = SendReceive,
                    li/LICn=ff00ff00ff00ff00},
                    li/LITID = [ffffff00, ffffff01],
                Remote {
                        v=0
                        c=IN IP4 124.124.124.222
                        m=audio 2222 RTP/AVP 0
                        a=ptime:20
                    }
                }
        }
    }
}
```

- The LICn command provides an encrypted correlation number. The PGCPD process decrypts the correlation number to determine whether session mirroring is performed on the gate. If the number is valid, interception is performed on the gate. If the number is invalid, no interception is performed on the gate.
- The LITID command contains one or more target IDs that identify the recipients of mirrored packets. If the request contains multiple LITIDs, a copy of mirrored

packets is created for each target ID. All copies are sent to the same delivery function. A maximum number of seven copies is supported for each packet.

### *How Session Mirroring Works*

If session mirroring is required on a gate, the PGCPD process embeds appropriate data in the gate open/modify request that it sends to the PIC. This data includes direction information to indicate whether the packet is mirrored before applying NAT actions or after. It also includes the decrypted correlation number and Target IDs that need to be embedded in the packet sent to the delivery function.

The PIC then:

1. Marks the gate that needs to be mirrored and obtains the destination for the mirrored packets from the CLI configuration.
2. Processes the packets as it normally does. It applies DSCP, latching, and rate limiting as appropriate.
   - Additional mirrored packets that are sent as a result of session mirroring do not impact rate limiting. The replicated packets do not count against policer counters that are used to compute the rate for the gate.
   - Mirrored packets are sent with DSCP marks that are applied to the gate as they are for a normal flow.
   - If the original packet is dropped because of rate limiting; no mirroring occurs.
   - Latch or relatch actions on the gate do not impact mirroring.
3. Generates one copy of the packets received on mirrored gates for each target ID specified in the PGCP request, encapsulates the mirrored packets, and sends them to the configured delivery function.

Session mirroring can be enabled or disabled any time during a gate's life by employing PGCP commands. If mirroring is enabled in one stream of a termination, all streams in the context are mirrored. Both RTP and RTCP packets are mirrored for a gate marked for mirroring.

### *Security for Packets Sent to the Delivery Function*

To protect mirrored traffic that is sent from the PG to the delivery function, you can use IPSec.

**Related Topics**
- Configuring Session Mirroring on page 78
- Displaying Gates That Are Being Mirrored on page 91

Chapter 4

# Configuring the Voice Solution

This chapter explains how to configure the voice solution. Topics include:



## Configuring a Virtual Packet Gateway

You can configure four virtual packet gateways (VPGs) on a router. Each VPG is associated with a different PIC.

**Step-by-Step Procedure** To configure a VPG:

1. Create a VPG, and assign a name to the VPG. You can configure an IP address as the VPG name. However, the IP address is not used in the operation of the VPG.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1
   ```

2. Specify the IP address of the VPG. This address is the local IP address on which the VPG receives Packet Gateway Control Protocol (PGCP) messages from the packet gateway controller (PGC).

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set gateway-address 10.10.30.1
   ```

3. Specify the port number of the VPG.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set gateway-port 2944
   ```

4. Configure the number of seconds before the VPG removes gates following a disconnection from the PGC.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set cleanup-timeout 3600
   ```

5. Configure the maximum number of concurrent calls allowed on the VPG. If you configure multiple VPGs for one service PIC, you can use this statement to achieve international oversubscription of resources or a fair distribution of resources between the VPGs.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set max-concurrent-calls 6000
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- Configuring the Virtual Packet Gateways on page 144

## Adding a PGC to the VPG Configuration

To configure a PGC for a VPG, specify the IP address and port number of the PGC. You can also secure the connection by specifying an algorithm for the interim AH scheme.

**Step-by-Step Procedure**

To configure a PGC for a VPG:

1. Access the configuration of the VPG for which you want to add a PGC

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1
   ```

2. Create a PGC configuration, and assign a name to the PGC. You can configure an IP address as the PGC name. However, the IP address is not used for the connection to the PGC.

```
[edit services pgcp gateway vpg-1]
user@host#edit gateway-controller pgc-1
```

3. Specify the IP address of the PGC.

```
[edit services pgcp gateway vpg-1 gateway-controller pgc-1]
user@host#set controller-address 10.10.2.3
```

4. Configure the number of the PGC listening port. The VPG sends H.248 messages to this port.

```
[edit services pgcp gateway vpg-1 gateway-controller pgc-1]
user@host#set controller-port 2944
```

5. To use the interim authentication header (AH) scheme to provide security on the PGCP connection, configure the security algorithm that the interim AH scheme uses. Currently, HMAC null is the only algorithm supported.

```
[edit services pgcp gateway vpg-1 gateway-controller pgc-1]
user@host#set interim-ah-scheme algorithm hmac-null
```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring NAT Pools for the Packet Gateway

You can configure a network address translation (NAT) pool that is exclusive to the packet gateway feature. This topic shows how to create the following types of NAT pools for PGCP:

- Configuring a Remotely Controlled NAT Pool on page 63
- Configuring a NAT Pool Selected Based on Transport Protocol on page 64

### *Configuring a Remotely Controlled NAT Pool*

**Step-by-Step Procedure** To configure a remotely-controlled NAT pool:

1. Create a NAT pool, and specify a name for the pool.

```
[edit services]
user@host#edit nat pool pgcp-pool
```

2. Configure an address range for the pool.

```
[edit services nat pool pgcp-pool]
user@host#set address-range low 10.10.20.100 high 10.10.30.100
```

3. Configure a range of ports. If you configure the NAT pool as remotely controlled, you must set a specific port range rather than using an automatic assignment of ports.

```
[edit services nat pool pgcp-pool]
user@host#set port range low 10000 high 50000
```

4. Specify that the NAT pool is used exclusively by the packet gateway.

```
[edit services nat pool pgcp-pool]
user@host#set pgcp
```

5. Specify that the PGC controls the addresses and ports in a NAT pool. The PGC reserves the addresses and ports when it requests specific local NAT bindings for remote addresses. (By default, the PG controls the addresses and ports in a pool.)

```
[edit services nat pool pgcp-pool]
user@host#set pgcp remotely-controlled
```

6. Configure the number of ports allocated to voice and video flows on the MultiServices PIC. This value is useful when one port is allocated for Real-Time Transport Protocol (RTP), and the accompanying Real-Time Control Protocol (RTCP) flow uses the other port. By default, 2 ports are available. To support the extra ports required for combined voice and video flows, you can specify 4 ports.

```
[edit services nat pool pgcp-pool]
user@host#set pgcp ports-per-session 4
```

### *Configuring a NAT Pool Selected Based on Transport Protocol*

**Step-by-Step Procedure** To configure a NAT pool that can be selected based on its transport protocol:

1. Create a NAT pool, and specify a name for the pool.

```
[edit services]
user@host#edit nat pool pgcp-pool
```

2. Configure an address range for the pool.

```
[edit services nat pool pgcp-pool]
user@host#set address-range low 10.10.20.100 high 10.10.30.100
```

3. Configure a range of ports. If you configure the NAT pool as remotely controlled, you must set a specific port range rather than using an automatic assignment of ports.

```
[edit services nat pool pgcp-pool]
user@host#set port range low 10000 high 50000
```

4. Specify that the NAT pool is used exclusively by the packet gateway.

```
[edit services nat pool pgcp-pool]
user@host#set pgcp
```

5. Specify one or more transport protocols that must match the transport protocol in the media descriptor of Add or Modify requests from the PGC.

```
[edit services nat pool pgcp-pool]
user@host#set pgcp transport [rtp-avp udp]
```

6. Specify that the port is automatically assigned.

```
[edit services nat pool pgcp-pool]
user@host#set port automatic
```

7. Configure the number of ports allocated to voice and video flows on the MultiServices PIC. This value is useful when one port is allocated for Real-Time Transport Protocol (RTP), and the accompanying Real-Time Control Protocol (RTCP) flow uses the other port. By default, 2 ports are available. To support the extra ports required for combined voice and video flows, you can specify 4 ports.

```
[edit services nat pool pgcp-pool]
user@host#set pgcp ports-per-session 4
```

**Related Topics**

- *Chapter 9, Summary of Network Address Translation Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- Twice NAT for VoIP Traffic Overview on page 49

## Assigning a NAT Pool

To assign a NAT pool, you create a media service configuration that contains the name of the NAT pool. You then specify the media service in a virtual interface configuration and in a PGCP rule. The PGCP rule assigns the media service for a specific VPG.

**Step-by-Step Procedure**

To configure a media service:

1. Create a media service, and specify a name for the service.

```
[edit services pgcp]
user@host#edit media-service media-service-one
```

2. Assign the NAT pool to the media service.

```
[edit services pgcp media-service media-service-one]
user@host#set nat-pool pgcp-pool
```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Virtual Interfaces

A voice call traverses two virtual interfaces. The configuration of the two virtual interfaces determines the NAT pool (media service) and the logical router interface (that is, a specific unit on a physical interface) for the call to use.

A virtual interface provides the mapping between interface names that appear in the H.248 termination ID in H.248 messages and the following:

- The media service to be used for a gate
- A physical interface on the router

**Step-by-Step Procedure** To configure a virtual interface:

1. Create a virtual interface, and specify a name for the interface.

   ```
   [edit services pgcp]
   edit virtual-interface v1
   ```

2. Specify the name of the media service that contains the NAT pool to be used for gates on the virtual interface that you are configuring.

   ```
   [edit services pgcp virtual-interface v1]
   user@host#set media-service media-service-one
   ```

3. Specify the physical router interface.

   ```
   [edit services pgcp virtual-interface v1]
   user@host#set interface fe-1/0/0
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Packet Gateway Rules

Packet gateway rules are applied on the MultiServices PIC. Packet gateway rules combined with firewall rules specify to the PIC how to deal with VoIP traffic. The packet gateway rules specify the NAT pool (media service) used on a specific VPG.

**Step-by-Step Procedure** To configure a packet gateway rule:

1. Create a rule and specify a name for the rule.

   ```
   [edit services pgcp]
   user@host#edit rule pgcp-rule-1
   ```

2. Specify the VPG on which this rule is applied.

   ```
   [edit services pgcp rule pgcp-rule-1]
   user@host#set gateway vpg-1
   ```

3. Specify the media service that contains the NAT pool to be used for this VPG.

   ```
   [edit services pgcp rule pgcp-rule-1]
   user@host#set media-service media-service-one
   ```

You can also configure packet gateway rules within a service set.

**Related Topics**
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring a Packet Gateway Rule Set

If you have defined multiple rules, you can specify the order in which the packet gateway processes the rules by creating a rule set. The packet gateway processes the rules in the order in which you specify them in the rule set. It processes rules as follows:

- If a rule matches the packet, the packet gateway performs the corresponding action and the rule processing stops.
- If no rule matches the packet, processing continues to the next rule in the set. If none of the rules match the packet, the packet is dropped by default.

**Step-by-Step Procedure** To configure a rule set.

1. Create a rule set and specify a name for the rule set.

```
[edit services pgcp]
user@host#edit rule-set pgcp-rule-set-1
```

2. Add a rule to the rule set.

```
[edit services pgcp rule-set pgcp-rule-set-1]
user@host#set rule pgcp-rule-1
```

3. Add additional rules to the rule set.

```
[edit services pgcp rule-set pgcp-rule-set-1]
user@host#set rule pgcp-rule-2
```

You can also configure packet gateway rule sets within a service set.

**Related Topics**
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring a Stateful Firewall for the Packet Gateway

**Step-by-Step Procedure** To create a stateful firewall:

1. Create a stateful firewall rule.

```
[edit services stateful-firewall]
user@host#edit rule r1
```

2. Set the match direction for the rule.

```
[edit services stateful-firewall rule r1]
user@host#set match-direction input-output
```

3. Add a term to the rule.

```
[edit services stateful-firewall rule r1]
user@host#set term t1 then reject
```

**Related Topics**

- *Chapter 7, Summary of Stateful Firewall Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring a Service Set

A service set allows you to combine directional rules, stateful firewall rules, CoS rules, and other rules that control the behavior of each service in the service set. We recommend you configure the PGCP service set as a next-hop service set that specifies the inside and outside logical service interfaces on the MultiServices PIC. You need to configure a service set for each PIC.

**Step-by-Step Procedure**

To configure a service set:

1. Create a service set configuration.

   ```
   [edit services]
   user@host#edit service-set pgcp-svc-set
   ```

2. Configure service set as a next-hop service set.

   ```
   [edit services service-set pgcp-svc-set]
   user@host#edit next-hop-service
   ```

3. Specify the service interface to the inside network. The interface must be a logical interface on the same service PIC and must not be used by another service set. This unit number must match the unit number of the inside service domain configured in the service interface.

   ```
   [edit services service-set pgcp-svc-set next-hop-service]
   user@host#set inside-service-interface sp-1/2/0.10
   ```

4. Specify the service interface to the outside network. The interface must be a logical interface on the same MultiServices PIC and must not be used by another service set. This unit number must match the unit number of the outside service domain configured in the service interface.

   ```
   [edit services service-set pgcp-svc-set next-hop-service]
   user@host#set outside-service-interface sp-1/2/0.20
   ```

5. Specify the name of the PGCP rule or rule set that applies to this service set.

   ```
   [edit services service-set pgcp-svc-set]
   user@host#set pgcp-rules pgcp-rule-1
   ```

6. Specify the name of the stateful firewall rule that applies to this service set.

   ```
   [edit services service-set pgcp-svc-set]
   user@host#set stateful-firewall-rules r1
   ```

7. Specify the name of the CoS rule that applies to this service set.

```
[edit services service-set pgcp-svc-set]
user@host#set cos-rules cos-rule
```

8. Configure logging for the service set.

```
[edit services service-set pgcp-svc-set]
user@host#edit syslog host local-1
[edit services service-set pgcp-svc-set syslog host local-1]
user@host#set services any
```

**Related Topics**

- *Chapter 22, Service Set Configuration Guidelines* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 23, Summary of Service Set Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Rate-Limiting for Voice Calls

**Step-by-Step Procedure** To configure rate-limiting for voice calls:

1. Access the configuration of the H.248 traffic-management properties.

```
[edit services pgcp gateway vpg-1]
user@host#edit h248-properties traffic-management
```

2. Configure a peak data rate for all gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set peak-data-rate-default 2000000
```

3. Configure a separate peak data rate for RTCP gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set peak-data-rate rtcp fixed-value 100000
```

4. Configure a sustained data rate for all gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set sustained-data-rate default 3000000
```

5. Configure a sustained data rate for RTCP gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set sustained-data-rate rtcp fixed-value 200000
```

6. Configure a maximum burst size for all gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set max-burst-size default 3000000
```

7. Configure a maximum burst size for RTCP gate streams.

```
[edit services pgcp gateway vpg-1 h248-properties traffic-management]
user@host#set max-burst-size rtcp percentage 1000
```

**Related Topics**
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring QoS for Voice Calls

**Step-by-Step Procedure** To configure a default DSCP value:

1. Access the configuration of the H.248 DiffServ properties.

```
[edit services pgcp gateway vpg-1]
user@host#edit h248-properties diffserv
```

2. Configure a value for the DSCP.

```
[edit services pgcp gateway vpg-1 h248-properties diffserv]
user@host#set dscp default ef
```

**Related Topics**
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring the Physical Interface to Advertise the VPG Address

You must configure one physical interface or one loopback interface on the router. On this interface, you configure the address of one or more VPGs. The physical or loopback interface advertises the VPG addresses through the routing protocols.

**Step-by-Step Procedure** To configure the router to advertise VPG addresses:

1. Configure the interface, and enter edit mode for the interface.

```
[edit interfaces]
user@host#edit fe-1/1/0 unit 0 family inet
```

2. Specify the IP address of each VPG. This address is the local IP address on which the VPG receives H.248 messages from the packet gateway controller PGC.

```
[edit interfaces fe-1/1/0 unit 0 family inet]
user@host#set address 10.10.30.1

[edit interfaces fe-1/1/0 unit 0 family inet]
user@host#set address 10.20.20.20
```

## Configuring the Service Interface

You need to configure a service interface that is a logical interface on the MultiServices PIC.

**Step-by-Step Procedure** To configure the service interface:

1. Configure the interface, and enter edit mode for the interface.

```
[edit interfaces]
user@host#edit sp-1/2/0
```

2. Configure a description for the interface.

```
[edit interfaces sp-1/2/0]
user@host#set description pgcp_service
```

3. Configure logical unit 0.

```
[edit interfaces sp-1/2/0]
user@host#set unit 0 family inet
```

4. Configure a logical unit and specify the protocol family.

```
[edit interfaces sp-1/2/0]
user@host#set unit 10 family inet
```

5. Set the service domain of the logical unit to inside. This unit number must match the unit number of the inside service interface configured in the service set.

```
[edit interfaces sp-1/2/0]
user@host#set unit 10 service-domain inside
```

6. Configure a logical unit and specify the protocol family.

```
[edit interfaces sp-1/2/0]
user@host#set unit 20 family inet
```

7. Set the service domain of the logical unit to outside. This unit number must match the unit number of the outside service interface configured in the service set.

```
[edit interfaces sp-1/2/0]
user@host#set unit 20 service-domain outside
```

8. Configure system logging on the service interface.

```
[edit interfaces sp-1/2/0]
user@host#set services-options syslog host local services any
```

9. Configure an inactivity timeout for sessions on the service interface.

```
[edit interfaces sp-1/2/0 services-options syslog]
user@host#set services-options inactivity-timeout 720
```

**Related Topics**

- *Chapter 24, Interface Configuration Guidelines* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 25, Summary of Interface Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring VPN Aggregation

VPN aggregation configurations have the following requirements.

- All interfaces in a pool must belong to the same service PIC.
- Logical interfaces cannot be in more than one pool.
- All interfaces must have either family inet or family inet6 configured.
- Logical unit 0 cannot be configured in a service interface pool.
- The maximum number of service interfaces in a pool is 1000.
- The service set must have a next-hop service that is set to the service interface pool; it cannot have inside and outside services.
- The service set must have a PGCP rule.
- In the virtual interface configuration, the service and physical interfaces must match interfaces configured in the VRF routing instance.
- The virtual interface configuration must include all media services in the PGCP rule that is configured in the service set.

**Step-by-Step Procedure**

See Figure 23 on page 58 for an illustration of this configuration.

To configure VPN Aggregation:

1. Configure a policy statement to be used for the vrf-import and vrf-export policies that you plan to configure in the routing instances.

```
[edit]
user@host#edit policy-options policy-statement policy-1

[edit policy-options policy-statement policy-1]
user@host#set term t1 then reject
```

2. Configure a VRF routing instance for each VPN.

```
[edit]
user@host#edit routing-instances vrf1

[edit routing-instances vrf1]
user@host#set instance-type vrf
user@host#set interface fe-1/3/0.10
user@host#set interface sp-1/0/0.10
user@host#set route-distinguisher 10.10.10.11:0
user@host#set vrf-import policy-1
user@host#set vrf-export policy-1

[edit]
user@host#edit routing-instances vrf2

[edit routing-instances vrf2]
user@host#set instance-type vrf
user@host#set interface ge-1/0/0.10
user@host#set interface sp-2/0/0.10
user@host#set route-distinguisher 10.10.10.22:0
```

```
user@host#set vrf-import policy-1
user@host#set vrf-export policy-1

[edit]
user@host#edit routing-instances vrf3

[edit routing-instances vrf1]
user@host#set instance-type vrf
user@host#set interface fe-1/2/0.10
user@host#set interface sp-3/0/0.10
user@host#set route-distinguisher 10.10.10.33:0
user@host#set vrf-import policy-1
user@host#set vrf-export policy-1
```

3. Configure a pool of logical PGCP service interfaces that are configured in the VRF routing instances.

```
[edit]
user@host#edit services service-interface-pools pool pgcp-pool

[edit services service-interface-pools pool pgcp-pool]
user@host#set interface sp-1/0/0.10
user@host#set interface sp-2/0/0.10
user@host#set interface sp-3/0/0.10
```

4. Create a service set that links the VRF and the PGCP services. Specify the service interface pool as the next-hop service. The service set must contain a PGCP rule. It cannot contain any other type of rule.

```
[edit]
user@host#edit services service-set pgcp

[edit services service-set pgcp]
user@host#set next-hop-service service-interface-pool pgcp-pool
user@host#set pgcp-rules pgcp-rule
```

5. Configure a virtual interface in the PGCP configuration for each VRF routing instance.

```
[edit]
user@host#edit services pgcp virtual-interface vif-1

[edit services pgcp virtual-interface 1]
user@host#set routing-instance vrf1 service-interface sp-1/0/0.10
user@host#set interface fe-1/3/0.10
user@host#set media-service access_ms

[edit]
user@host#edit services pgcp virtual-interface vif-2

[edit services pgcp virtual-interface 2]
user@host#set routing-instance vrf1 service-interface sp-2/0/0.10
user@host#set interface ge-1/0/0.10
user@host#set media-service core_ms

[edit]
user@host#edit services pgcp virtual-interface vif-3
```

```
[edit services pgcp virtual-interface 1]
user@host#set routing-instance vrf3 service-interface sp-3/0/0.10
user@host#set interface fe-1/2/0.10
user@host#set media-service access_ms
```

**Related Topics**

- VPN Aggregation for VoIP Calls Overview on page 56
- *JUNOS VPNs Configuration Guide*
- *Chapter 23, Summary of Service Set Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 29, Summary of Service Interface Pools Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Latch Deadlock and Media Inactivity Detection

The PGC can request to be notified by the VPG when a latching deadlock or media inactivity exists on a PGCP gate. You can configure latching deadlock and media inactivity detection parameters for use by the VPG when it monitors media traffic that is flowing through the gate.

Latch deadlock detection is defined in *Gateway Control Protocol: Application Data Inactivity Package, ITU-T Recommendation H.248.40, January, 2007.*

**Step-by-Step Procedure** To configure parameters for latching deadlock and media inactivity detection:

1. Access the configuration of your VPG and specify data-inactivity-detection.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1 data-inactivity-detection
   ```

2. Configure the number of seconds before the VPG begins checking for media inactivity on new gates for which there is a latching signal.

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#set latch-deadlock-delay 10
   ```

3. Configure the number of seconds before the VPG begins checking for media inactivity on new gates that do not have a latching signal.

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#set inactivity-delay 10
   ```

4. Configure the duration of inactivity detection checks that the VPG performs on a gate. If no media packets are received during a check, the VPG sends the PGC a quality (QUA) alert, ADID alert, or service change notification. If media packets are received, the timer is reset and checking continues. This parameter applies to all gates, regardless of whether there is a latching signal for the gate.

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#set inactivity-duration 60
   ```

5. Request a service change to take gates with latch deadlocks or media inactivity out of service, dropping all packets for the gates. Specify whether to notify the PGC with error code 906 (loss of lower layer connectivity) or 910 (media capability failure).

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#edit report-service-change

   [edit services pgcp gateway vpg-1 data-inactivity-detection report-service-change]
   user@host#set service-change-type forced-910
   ```

6. Configure the VPG to stop inactivity detection when a gate action is set to drop. Use this option to handle calls that are placed on hold. When calls are resumed, the PG starts the delay timer and resumes data inactivity detection. By default, inactivity detection continues when a gate option is ready to drop

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#set stop-detection-on-drop
   ```

7. Specify that a notification (or service change) occur immediately when no media packets are detected during the initial checking delay period (latch-deadlock-delay or inactivity-delay). By default, inactivity is not reported until the delay period and an inactivity duration period have elapsed.

   ```
   [edit services pgcp gateway vpg-1 data-inactivity-detection]
   user@host#set send-notification-on-delay
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring H.248 Timers

You can configure H.248 timers for the PGCP connection between the VPG and the PGC. See clause 9.2 and annex D of the *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*, for details about these timers.

**Step-by-Step Procedure**

To configure H.248 timers for the PGCP connection between the VPG and the PGC:

1. Access the configuration of the H.248 timers.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1 h248-timers
   ```

2. Configure the value of the average acknowledgment delay (AAD) that the VPG uses before the first AAD is measured.

   ```
   [edit services pgcp gateway vpg-1 h248-timers]
   user@host#set initial-average-ack-delay 1000
   ```

3. Configure the assumed maximum network propagation delay time.

   ```
   [edit services pgcp gateway vpg-1 h248-timers]
   user@host#set maximum-net-propagation-delay 5000
   ```

4. Configure a maximum waiting delay (MWD) that is used when the VPG attempts to reconnect to the PGC. If the VPG finishes traversing its list of PGCs, and has not connected to a PGC, the VPG waits for a random value between 0 and MWD milliseconds before it attempts to reconnect to a PGC.

   ```
   [edit services pgcp gateway vpg-1 h248-timers]
   user@host#set maximum-waiting-delay 10000
   ```

5. Configure the maximum time that a transaction can be kept alive. When this time expires, the packet gateway considers the PGC to be down.

   ```
   [edit services pgcp gateway vpg-1 h248-timers]
   user@host#set tmax-retransmission-delay 25000
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Default Values for H.248 Base Root Properties

You can configure default values for properties defined in the H.248 base root package. The VPG uses these values unless the PGC overrides them with a PGCP command. In general we recommend you use the values that are configured on the router by default.

The base root package is defined in annex E.2 of the *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*. The properties in this package mostly affect the timers used when the VPG and PGC send and receive provisional responses to H.248 commands.

**Step-by-Step Procedure**

To configure H.248 timers for the PGCP connection between the VPG and the PGC:

1. Access the configuration of the H.248 base root properties.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1 h248-properties base-root
   ```

2. Set the default value for the number of milliseconds for the PGC to wait for a response to transactions from the VPG.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set normal-mg-execution-time default 500
   ```

3. Set the default value for the number of milliseconds for the PGC to wait for a pending response from the VPG if a transaction cannot be completed.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set mg-provisional-response-timer-value default 2000
   ```

4. Set the default value for the number of transaction pending messages that the PGC can receive from the VPG.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set mg-originated-pending-limit default 4
   ```

5. Set the default value for the number of milliseconds for the VPG to wait for a response to transactions from the PGC.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set normal-mgc-execution-time default 500
   ```

6. Set the default value for the number of milliseconds for the VPG to wait for a pending response from the PGC if a transaction cannot be completed.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set mgc-provisional-response-timer-value default 4000
   ```

7. Set the default value for the number of transaction pending messages that the VPG can receive from the PGC.

   ```
   [edit services pgcp gateway vpg-1 h248-properties base-root]
   user@host#set mgc-originated-pending-limit default 4
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Default Values for H.248 Segmentation Properties

You can configure default values for properties defined in the H.248 segmentation package. The VPG uses these values unless the PGC overrides them with a PGCP command. In general, we recommend you use the values that are configured on the router by default.

The segmentation package is defined in annex E.14 of the *Gateway control protocol v3, ITU T Recommendation H.248.1, September 2005*. The properties in this package affect the limits used when long H.248 replies are segmented into several H.248 messages.

**Step-by-Step Procedure**

To configure default values for H.248 segmentation properties:

1. Access the configuration of the H.248 segmentation properties.

   ```
   [edit services pgcp ]
   user@host#edit gateway vpg-1 h248-properties segmentation
   ```

2. Set a default value for the number of milliseconds for the PGC to wait for outstanding message segments from the VPG after it receives the SegmentationCompleteToken message.

   ```
   [edit services pgcp gateway vpg-1 h248-properties segmentation]
   user@host#set mg-segmentation-timer default 4000
   ```

3. Set a default value, in bytes, for the MG maximum PDU size property of the segmentation package. This value determines the maximum size of messages that the PGC sends to the packet gateway.

   ```
   [edit services pgcp gateway vpg-1 h248-properties segmentation]
   user@host#set mg-maximum-pdu-size default 1472
   ```

4. Set a default value for the number of milliseconds for the VPG to wait for outstanding message segments from the PGC after it receives the SegmentationCompleteToken.

   ```
   [edit services pgcp gateway vpg-1 h248-properties segmentation]
   user@host#set mgc-segmentation-timer default 4000
   ```

5. Set a default value, in bytes, for the MGC maximum PDU size property of the segmentation package. This value determines the maximum size of messages that the VPG sends to the PGC.

   ```
   [edit services pgcp gateway vpg-1 h248-properties segmentation]
   user@host#set mgc-maximum-pdu-size default 1472
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Enabling Wildcards for Service Change Notifications

You can enable the VPG to issue service change commands as wildcard-response commands, which trigger a short response from the PGC. If you do not enable the use of wildcard response for service change commands, the PGC generates an individual response for every termination that matches the service change command.

**Step-by-Step Procedure**

To enable wildcard-response commands for service change commands:

1. Access the H.248 options configuration.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#edit h248-options
   ```

2. Enable the VPG to issue service change commands as wildcard-response commands.

   ```
   [edit services pgcp gateway vpg-1 h248-options]
   user@host#set wildcard-response-service-change
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Configuring Session Mirroring

Session mirroring commands are hidden by default. You must have a login with sufficient permission to configure session mirroring. The set system login class *class-name* permissions pgcp-session-mirroring-control command grants this permission.

**Step-by-Step Procedure**

To configure session mirroring:

1. Access the configuration of the delivery function properties under session-mirroring.

   ```
   [edit services pgcp ]
   ```

```
user@host#edit session-mirroring delivery-function df-1
```

2. Configure the network operator ID. The PG includes the network operator ID in the header of intercepted packets that it sends to the delivery function. It is used to identify the operator.

```
[edit services pgcp session-mirroring delivery-function df-1]
user@host#set network-operator-id ABCDE
```

3. Configure the address of the delivery function to which the PG sends session-mirroring information.

```
[edit services pgcp session-mirroring delivery-function df-1]
user@host#set destination-address 10.1.1.63
```

4. Configure the port on the delivery function that receives session-mirroring information.

```
[edit services pgcp session-mirroring delivery-function df-1]
user@host#set destination-port 15000
```

5. Configure the address of the interface on which the PG sends session-mirroring data to the deliver function.

```
[edit services pgcp session-mirroring delivery-function df-1]
user@host#set source-address 10.1.1.43
```

6. Configure the port on which the PG sends session-mirroring data to the delivery function.

```
[edit services pgcp session-mirroring delivery-function df-1]
user@host#set source-port 10000
```

### *Disabling Session Mirroring*

To disable session mirroring:

```
[edit services pgcp session-mirroring]
user@host#set disable-session-mirroring
```

### *Re-Enabling Session Mirroring*

To re-enable session mirroring:

```
[edit services pgcp session-mirroring]
user@host#delete disable-session-mirroring
```

### *Configuring IPSec for Mirrored Sessions*

To protect mirrored traffic that is sent from the PG to the delivery function, you can use IPSec. To have IPSec and PGCP performed on the same PIC, you create PGCP and IPSec service sets and chain these service-sets using routing-options.

To create the service sets and routing options:

1. Configure a PGCP service set. The NAT routes installed as part of PGCP service direct PGCP traffic to sp-1/0/0.10 and sp-1/0/0.20.

   ```
   [edit services service-set pgcp-svc-set]
   user@host#set pgcp-rules pgcp-rule
   user@host#set next-hop-service inside-service-interface sp-1/0/0.10
   user@host#set next-hop-service outside-service-interface sp-1/0/0.20
   ```

2. Configure an IPSec service set on the same PIC.

   ```
   [edit services service-set ipsec-svc-set]
   user@host#set next-hop-service inside-service-interface sp-1/0/0.30
   user@host#set next-hop-service outside-service-interface sp-1/0/0.40
   user@host#set ipsec-vpn-options local-gateway 1.0.0.1
   user@host#set ipsec-vpn-rules ipsec1
   ```

3. Install a static route to the delivery function (1.0.0.3) with the next-hop address of the PIC. This route redirects mirrored packets to a unit of the same service PIC that is hosting the IPSec service.

   ```
   [edit]
   user@host#set routing-options static route 1.0.0.3/32 next-hop sp-1/0/0.30
   ```

The mirrored packets that are generated on sp-1/0/0 have the destination address of the delivery function. In this case 1.0.0.3.

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Verifying Your Configuration

You can use show commands to verify your configuration.



### *Verifying the PGCP Configuration*

**Purpose** Display the active configuration.

**Action**

```
[edit services pgcp]
user@host# show
media-service media-service-one {
    nat-pool pgcp-pool;
}
```

```
virtual-interface v1 {
    media-service media-service-one;
    interface fe-1/0/0.0;
}
gateway vpg-1 {
    gateway-address 10.10.30.1;
    gateway-port 2944;
    cleanup-timeout 3600;
    gate-inactivity-delay 10;
    gate-inactivity-duration 60;
    h248-timers {
        maximum-waiting-delay 10000;
        tmax-retransmission-delay 25000;
        initial-average-ack-delay 1000;
        maximum-net-propagation-delay 5000;
    }
    h248-properties {
        base-root {
            normal-mg-execution-time {
                default 500;
            }
            mg-provisional-response-timer-value {
                default 2000;
            }
            mg-originated-pending-limit {
                default 4;
            }
            normal-mgc-execution-time {
                default 500;
            }
            mgc-provisional-response-timer-value {
                default 4000;
            }
            mgc-originated-pending-limit {
                default 4;
            }
        }
        segmentation {
            mgc-segmentation-timer {
                default 4000;
            }
            mgc-maximum-pdu-size {
                default 1472;
            }
            mg-segmentation-timer {
                default 4000;
            }
            mg-maximum-pdu-size {
                default 1472;
            }
        }
        diffserv {
            dscp {
                default ef;
            }
        }
        traffic-management {
            sustained-data-rate {
                default 3000000;
                rtcp {
                    fixed-value 200000;
```

```
                }
            }
            peak-data-rate {
                default 2000000;
                rtcp {
                    fixed-value 100000;
                }
            }
            max-burst-size {
                default 3000000;
                rtcp {
                    percentage 1000;
                }
            }
        }
    }
    h248-options {
        wildcard-response-service-change;
    }
    gateway-controller pgc-1 {
        controller-address 10.10.2.3;
        controller-port 2944;
        interim-ah-scheme {
            algorithm hmac-null;
        }
    }
}
rule pgcp-rule-1 {
    gateway vpg-1;
    media-service media-service-one;
}
rule-set pgcp-rule-set-1 {
    rule pgcp-rule-1;
}
session-mirroring {
    delivery-function df-1 {
        destination-address 10.1.1.63;
        destination-port 15000;
        network-operator-id ABCDE;
        source-address 10.1.1.43;
        source-port 10000;
    }
}
```

### *Verifying the Service Interface Configuration*

**Purpose** Display the service interface configuration.

**Action**

```
[edit interface sp-1/0/0]
user@host# show
description pgcp_service;
traceoptions {
    flag all;
}
services-options {
    syslog {
        host local {
            services any;
```

```
        }
    }
    inactivity-timeout 720;
}
unit 0 {
    family inet;
}
unit 10 {
    family inet;
    service-domain inside;
}
unit 20 {
    family inet;
    service-domain outside;
}
```

### *Verifying the Physical Interface Configuration*

**Purpose** Display the physical interface configuration.

**Action**
```
[edit interfaces fe-1/1/0]
user@host# show
unit 0 {
    family inet {
        address 10.10.30.1/32;
        address 10.20.20.20/32;
    }
}
```

### *Verifying the Service Set Configuration*

**Purpose** Display the service set configuration.

**Action**
```
[edit services service-set pgcp-svc-set]
user@host# show
syslog {
    host local-1 {
        services any;
    }
}
stateful-firewall-rules r1;
cos-rules cos-rule;
pgcp-rules pgcp-rule-1;
next-hop-service {
    inside-service-interface sp-1/2/0.10;
    outside-service-interface sp-1/2/0.20;
```

### *Verifying the NAT Pool Configuration*

**Purpose** Display the NAT pool configuration.

**Action**
```
[edit services nat]
user@host# show
pool pgcp-pool {
    pgcp {
        remotely-controlled;
```

```
            ports-per-session 4;
            }
        address-range low 10.10.20.100 high 10.10.30.100;
        port range low 10000 high 50000;
    }
```

### *Verifying the Stateful Firewall Configuration*

**Purpose** Display the stateful firewall configuration.

**Action**

```
[edit services stateful-firewall rule r1]
user@host# show
match-direction input-output;
term t1 {
    then {
        reject;
    }
}
```

Chapter 5

# Monitoring the Voice Solution

This chapter explains how to monitor the voice solution components. Topics include:



## Monitoring RTP and RTCP Traffic

To monitor Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) packets on media gates, the VPG uses RTP and RTCP application layer gateway (ALGs) that are attached to flows when the PGC installs media gates on the VPG.

The RTP ALG uses the sequence number of each RTP packet on that flow to record the number of lost packets. Each RTP packet contains a 16-bit sequence number field that is incremented for every packet sent. The starting sequence number is randomly selected.

The RTCP ALG monitors the Sender Report and Receiver Report packet types. For the sender, the ALG records the Synchronization Source (SSRC) value and the number of invalid packets, sender packets, and sender octets.

For the receiver, the RTCP ALG records monitors lost packets, the fraction of lost packets, and jitter. When the ALG tracks multiple receiver reports (that is, when a sender is listening to media from multiple sources), it tracks the statistics for up to four sources. When more than four sources are available, the ALG overwrites the statistics for the source with the oldest record.

### *Enabling Monitoring of RTP and RTCP Traffic*

You can enable RTP and RTCP ALGs for twice NAT flows created when the PGC installs media gates on the VPG. The ALGs monitor packets on the gate and provide statistics.

You can enable these ALGs only for flows created by the VPG. You cannot enable them within standalone NAT rules.

**Step-by-Step Procedure**

You can choose to monitor either RTP or RTCP, or both. To enable monitoring of RTP and RTCP media flows:

1. Access the configuration of your VPG.

```
[edit services pgcp]
user@host#edit gateway vpg-1
```

2. Enable monitoring of both RTP and RTCP for media flows.

```
[edit services pgcp gateway vpg-1]
user@host#set monitor media
```

3. Enable monitoring of only RTP media flows.

```
[edit services pgcp gateway vpg-1]
user@host#set monitor media rtp
```

4. Enable monitoring of only RTCP media flows.

```
[edit services pgcp gateway vpg-1]
user@host#set monitor media rtcp
```

5. In operational mode, view statistics gathered for RTP and RTCP.

```
user@host> show services pgcp gates gate-id 352187384065 statistics

Gate Statistics:
================
Output packets: 582
Input packets: 582
Dropped packets: 0
Lost RTP packets: 0

RTCP statistics:

SSRC           : 32270
Sender octets  : 7500
Sender packets : 375
Invalid packets: 9

RTCP Receiver statistics:

SSRC        Lost packets   Lost fraction         Jitter
13043                  0           0.000              0
16487                  0           0.000              0
5655                   0           0.000              0

. . .
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Monitoring Gates

Use the following topics to learn about monitoring gates:



### *Displaying Information About All Gates on a VPG*

**Purpose** Display information about all gates on a VPG using the show services pgcp gates gateway *gateway-name* command.

**Action**

```
user@host> show services pgcp gates gateway vpg-1
Packet gateway configuration:
  Name                                    :  vpg-1
  IP address                              :  3.0.0.2
  Port                                    :  2944
  Status                                  :  Connected

Gate information:
Gate id: 4295033088
Gate state: Active
Service set id: 1
Media card: sp-0/3/0
Media handler: pgcp-svc-set-1
Termination-id-string: ip/0/rlmvi2/1

Gate id: 4295033089
Gate state: Active
Service set id: 1
Media card: sp-0/3/0
Media handler: pgcp-svc-set-1
Termination-id-string: ip/0/rlmvi0/2

Gate id: 8590000384
Gate state: Active
Service set id: 1
Media card: sp-0/3/0
Media handler: pgcp-svc-set-1
Termination-id-string: ip/0/rlmvi2/3

Gate id: 8590000385
Gate state: Active
Service set id: 1
Media card: sp-0/3/0
Media handler: pgcp-svc-set-1
Termination-id-string: ip/0/rlmvi0/4
```

### *Displaying Extensive Information About All Gates on a VPG*

**Purpose** Display information about all gates on a VPG including statistics for the gates, RTCP sender and RTCP receiver statistics, and rate-limiting statistics for the gates, using the show services pgcp gates gateway *gateway-name* extensive command.

**Action**

```
user@host> show services pgcp gates gateway vpg-1 extensive
Packet gateway configuration:
  Name                                    :  vpg-1
  IP address                              :  2.0.0.2
  Port                                    :  2944
  Status                                  :  In-Service (Registered)

Gate information:
================
Gate id: 4295033088
Gate state: active
Direction: A->B
Action: forward
Remote source address: *
Remote source port: *
Remote destination address: 4.0.0.1
Remote destination port: 5060
Local source address: [4.99.99.20]
Local source port: [5060]
Local destination address: 2.99.99.20
Local destination port: 5060
Transport: udp
RTCP: Off
Latch: none
DSCP: 0x00 (Effective 0)
Policing: On
Gate SDR : 10000 bytes per second
Gate PDR : 10000 bytes per second
Gate MBS : 1000 bytes
RTCP SDR : 500 bytes per second
RTCP PDR : 500 bytes per second
RTCP MBS : 1000 bytes
Fast update filter: Off

Gate information:
================
Gate id: 4295033089
Gate state: active
Direction: B->A
Action: forward
Remote source address: *
Remote source port: *
Remote destination address: 2.0.0.1
Remote destination port: 5060
Local source address: [2.99.99.20]
Local source port: [5060]
Local destination address: 4.99.99.20
Local destination port: 5060
Transport: udp
RTCP: Off
Latch: none
DSCP: 0x00 (Effective 0)
```

```
Policing: On
Gate SDR : 10000 bytes per second
Gate PDR : 10000 bytes per second
Gate MBS : 1000 bytes
RTCP SDR : 500 bytes per second
RTCP PDR : 500 bytes per second
RTCP MBS : 1000 bytes
Fast update filter: Off
```

### *Displaying the Number of Gates Installed on a VPG*

**Purpose** Display the number of gates installed on a VPG using the show services pgcp gates gateway *gateway-name* count command.

**Action**

```
user@host> show services pgcp gates gateway vpg-1 count
Gateway Name                                        Gate count
vpg-1                                                      4
```

### *Displaying Information About a Specific Gate*

**Purpose** Display information about a specific gate using the show services pgcp gates gate-id *gate-id* command.

**Action**

```
user@host> show services pgcp gates gate-id 4295033089

Gate information:
Gate id: 4295033089
Gate state:  active
Action: forward
Service set id: 1
Media card: sp-0/2/0
Media handler: pgcp-svc-set
Termination-id-string: ip/0/rlmvi2/1
```

### *Displaying Extensive Information About a Specific Gate*

**Purpose** Display extensive information about a specific gate including statistics for the gate using the show services pgcp gates gate-id *gate-id* extensive command.

**Action**

```
user@host> show services pgcp gates gate-id 4295033089 extensive
Gate information:
================
Gate id: 4295033089
Gate state: active
Direction: B->A
Action: forward
Remote source address: *
Remote source port: *
Remote destination address: 2.0.0.1
Remote destination port: 5060
Local source address: [2.99.99.20]
Local source port: [5060]
Local destination address: 4.99.99.20
Local destination port: 5060
Transport: udp
```

```
RTCP: Off
Latch: none
DSCP: 0x00 (Effective 0)
Policing: On
Gate SDR : 10000 bytes per second
Gate PDR : 10000 bytes per second
Gate MBS : 1000 bytes
RTCP SDR : 500 bytes per second
RTCP PDR : 500 bytes per second
RTCP MBS : 1000 bytes
Fast update filter: Off
```

### *Displaying Statistics for Gates*

**Purpose** Display statistics for a gate, including packet statistics, RTCP sender and RTCP receiver statistics, rate-limiting statistics, and the number of packets dropped because of fast update filters (FUF), using the show services pgcp gates gate-id *gate-id* statistics command.

**Action**

```
user@host> show services pgcp gates gate-id 98784313601 statistics
Gate Statistics:
================
Output packets: 0
Input packets: 0
Dropped packets: 0
Lost RTP packets: 0

Rate limiting statistics:

Mark Color   Number of Packets    Number of Bytes

Green                        0                  0
Yellow                       0                  0
Red                          0                  0


FUF statistics:
Drop count: 0
```

### *Collecting Statistics on Gates with Rate-Limited Flows*

When fast update filters (FUF) are installed on the Packet Forwarding Engine and PIC, the Packet Forwarding Engine and PIC discard packets flowing through gates if they exceed the rate limits set in the FUF. To get accurate statistics for gates, the JUNOS software includes these discarded packets in statistics that it collects for gates.

### *Improving Performance While Collecting Gate Statistics*

Collecting statistics on packets that are dropped on a gate can impact system performance. To improve performance, the software can limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG. This limit is the maximum value of the following parameters configured on the router:

- The maximum number of FUF terms installed for the VPG

- The maximum percentage of gates with FUF filters relative to all gates currently installed for the VPG

**Step-by-Step Procedure** To configure a limit on the number of FUF terms installed on the Packet Forwarding Engine for a VPG:

1. Access the configuration of your VPG.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1
   ```

2. Specify the maximum number of FUF terms installed for the VPG.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set fast-update-filter maximum-terms 3000
   ```

3. Specify the maximum percentage of gates with FUF filters relative to all gates currently installed for the VPG.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#set fast-update-filter maximum-fuf-percentage 15
   ```

### Displaying the Number of FUF Terms Installed on a VPG

**Purpose** Display the number of match condition FUF terms and the number of FUF filters currently installed on a VPG using the show services pgcp active-configuration command.

**Action**

```
user@host> show services pgcp active-configuration
. . .
Firewall:
  Status             :  Connected
  Number of terms    :  2
  Number of filters  :  2
```

## *Displaying Gates That Are Being Mirrored*

You can view session mirroring information for gates that are being mirrored. You must have a login with sufficient permission to view session mirroring information. The set system login class *class-name* permissions pgcp-session-mirroring command grants this permission.

**Purpose** Display session-mirroring information for a gate using the show services pgcp gates gate-id *gate-id* session-mirroring command.

**Action**

```
user@host> show services pgcp gates gate-id 4295033088 session-mirroring

Gate information:
Gate id: 4295033088
Session mirroring status: On
Session mirroring correlation number: 0x8040c020a060e010
Session mirroring target ID list: [008040c0, ffffff80]
Session mirroring direction: Egress
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Monitoring PGCP Terminations

To monitor PGCP terminations, you can display the following information on the router:



### *Displaying Information About All PGCP Terminations on a VPG*

**Purpose** Display information about all PGCP terminations on a VPG using the show services pgcp terminations *gateway-name* command.

**Action**

```
user@host> show services pgcp terminations vpg-1
Packet gateway configuration:
  Name                                  :  vpg-1
  IP address                            :  2.0.0.2
  Port                                  :  2944
  Status                                :  In-Service (Registered)


Termination name                  State          Duration(msecs)
ip/4/vif-0/3                 In-service      920610
Gate-id         Direction       State        Action
4295033088      A->B            active       forward
4295033089      B->A            active       forward


Termination name                  State          Duration(msecs)
ip/4/vif-0/2         In-service      920618
Gate-id         Direction       State        Action
4295033088      A->B            active       forward
4295033089      B->A            active       forward
```

### *Displaying Information About PGCP Terminations in H.248 Format*

**Purpose** Display information about PGCP terminations in H.248 format using the show services pgcp terminations *gateway-name* h248 command.

**Action**

```
user@host> show services pgcp terminations vpg-1 h248
Termination information:
ip/4/vif-0/2 {

      MEDIA {

        TERMINATIONSTATE { SERVICESTATES = INSERVICE },

        STREAM = 1 {
```

```
          LOCALCONTROL { MODE = SENDRECEIVE,
                         DS/DSCP = 00,
                         TMAN/MBS = 5000,
                         TMAN/PDR = 0,
                         TMAN/POL = ON,
                         TMAN/SDR = 125000,
                         MGCINFO/DB = 00,
                         GM/RSB = OFF,
                         GM/SAF = OFF,
                         GM/SPF = OFF,
                         GM/SPR = 0,
                         GM/ESAS = OFF,
                         GM/ESPS = OFF,
                         GM/LSP = 0 },
          LOCAL {
v=0
c=IN IP4 4.99.99.20
m=- 5060 udp -
b=AS:0

          },
          REMOTE {
v=0
c=IN IP4 4.0.0.1
m=- 5060 udp -
b=AS:0

          }
        }
      },
      EVENTS { HANGTERM/THB { TIMERX= 30 } }
}
```

```
ip/4/vif-0/3 {
      MEDIA {
        TERMINATIONSTATE { SERVICESTATES = INSERVICE },
        STREAM = 1 {
          LOCALCONTROL { MODE = SENDRECEIVE,
                         DS/DSCP = 00,
                         TMAN/MBS = 5000,
                         TMAN/PDR = 0,
                         TMAN/POL = ON,
                         TMAN/SDR = 125000,
                         MGCINFO/DB = 00,
                         GM/RSB = OFF,
                         GM/SAF = OFF,
                         GM/SPF = OFF,
                         GM/SPR = 0,
                         GM/ESAS = OFF,
                         GM/ESPS = OFF,
                         GM/LSP = 0 },
          LOCAL {
v=0
c=IN IP4 2.99.99.20
m=- 5060 udp -
b=AS:0

          },
          REMOTE {
v=0
c=IN IP4 2.0.0.1
m=- 5060 udp -
b=AS:0
```

```
        }
      }
    },
    EVENTS { HANGTERM/THB { TIMERX= 30 } }

}
```

### *Displaying Information About Specific PGCP Terminations*

**Purpose** Display information about a specific PGCP termination using the show services pgcp terminations termination-prefix *termination-prefix* h248 *gateway-name* command. For the termination prefix, you can enter a partial name, and all matching terminations are displayed.

**Action**

```
user@host> show services pgcp terminations termination-prefix ip/4/vif-0/3 h248
vpg-1
Termination information:
ip/4/vif-0/3 {
    MEDIA {
      TERMINATIONSTATE { SERVICESTATES = INSERVICE },
      STREAM = 1 {
        LOCALCONTROL { MODE = SENDRECEIVE,
                       DS/DSCP = 00,
                       TMAN/MBS = 5000,
                       TMAN/PDR = 0,
                       TMAN/POL = ON,
                       TMAN/SDR = 125000,
                       MGCINFO/DB = 00,
                       GM/RSB = OFF,
                       GM/SAF = OFF,
                       GM/SPF = OFF,
                       GM/SPR = 0,
                       GM/ESAS = OFF,
                       GM/ESPS = OFF,
                       GM/LSP = 0 },
        LOCAL {
```

```
v=0

c=IN IP4 4.99.99.20

m=- 5060 udp -

b=AS:0


          },

          REMOTE {

v=0

c=IN IP4 4.0.0.1

m=- 5060 udp -

b=AS:0


          }

        }

      },

      EVENTS { HANGTERM/THB { TIMERX= 30 } }

}
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Monitoring PGCP Root Terminations

**Purpose** Display information about the root termination on a VPG using the show services pgcp root-termination *gateway-name* command.

**Action**

```
user@host> show services pgcp root-termination vpg-1
Root termination information:
ROOT {

      MEDIA {

        TERMINATIONSTATE { SERVICESTATES = OUTOFSERVICE,

                           ROOT/MAXNUMBEROFCONTEXTS = 21000,

                           ROOT/MAXTERMINATIONSPERCONTEXT = 2,

                           ROOT/MGCORIGINATEDPENDINGLIMIT = 4,

                           ROOT/MGCPROVISIONALRESPONSETIMERVALUE = 4000,

                           ROOT/MGORIGINATEDPENDINGLIMIT = 4,
```

```
                                ROOT/MGPROVISIONALRESPONSETIMERVALUE = 2000,
                                ROOT/NORMALMGCEXECUTIONTIME = 500,
                                ROOT/NORMALMGEXECUTIONTIME = 500,
                                SEG/MGCMAXPDUSIZE = 1472,
                                SEG/MGCSEGMENTATIONTIMERVALUE = 4000,
                                SEG/MGMAXPDUSIZE = 1472,
                                SEG/MGSEGMENTATIONTIMERVALUE = 4000 }
      },
      EVENTS = 3 { IT/ITO { MIT = 12000 } }
    }
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Monitoring Statistics for PGCP

**Purpose** Display statistics for H.248 messages, protocol errors, and commands using the show services pgcp statistics gateway *gateway-name* command.

**Action**

```
user@host> show services pgcp statistics gateway vpg-1
Packet gateway configuration:
  Name                                  :  vpg-1
  IP address                            :  3.0.0.2
  Port                                  :  2944
  Status                                :  Connected

H.248 statistics:
    Messages received          :  303852
    Messages sent              :  303726
    Protocol errors            :  329

Command Requests               sent/wildcards            received/wildcards
    Add                             0/0                       207557/0
    Modify                          0/0                       206955/0
    Notify                         33/0                            0/0
    ServiceChange                 330/0                            0/0
    Subtract                        0/0                       191195/0

Command Responses              received/error               sent/error
    Add                             0/0                       207557/378
    Modify                          0/0                       206898/110
    Notify                         33/0                            0/0
    ServiceChange                   1/0                            0/0
    Subtract                        0/0                       191073/503
```

**Meaning** There is no relationship between the messages received and sent counters and the command request and response counters. The message counters describe the network

activity and contain all H.248 messages. Each message can contain one or more commands, and commands that contain wildcards can be interpreted as more than one command.

### *Command Requests*

The command requests section lists the number of command requests that the VPG sent and received. It lists the number of requests for each type of command. Command requests are counted regardless of whether the commands were successfully accomplished. Command requests are not counted in the following cases:

- The command was not executed because of a previous error.
- The command was not fully executed because of its own syntax error, which made it impossible to obtain the command type itself.

The number of requests are divided into sent and received requests as follows:

- Sent and received—Commands that have a specific context ID and termination ID.
- Wildcards—Commands that contain wildcard context IDs or wildcard termination IDs or both. The counter counts the total number of matches on wildcards.

### *Command Responses*

Command responses received do not reflect the number of matches sent in command responses because:

- Wildcard responses do not indicate how many matches succeeded
- The software stops parsing responses when it first encounters an error so there is no distinction between single and wildcard commands.

Command responses received are incremented as follows:

- Each response received increments the received counter.
- Each command failure increments the error counter.

The command responses sent/error counters are incremented as follows:

- Each successful match increments the sent counter.
- Each failure increments both the sent and the error counter.

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Monitoring PGCP Flows

You can view PGCP flows with the show services pgcp flows command. You can view information for all PGCP flows, or you can include keywords to limit the number of flows displayed or to display flows for a particular:

- VPG
- Gate
- Application protocol
- Source or destination port
- Source or destination prefix
- Protocol
- Service set

### *Displaying All PGCP Flows*

**Purpose** Display standard information about all PGCP flows using the show services pgcp flows command.

**Action**

```
user@host> show services pgcp flows
Interface: sp-0/3/0, Service set: pgcp-svc-set-1
Flow                                                State    Dir       Frm count
UDP          4.0.0.102:0     ->    4.99.99.100:1024  Forward  I           21531
Gate id: 8590000385
    NAT source       4.0.0.102:0       ->     3.99.99.100:1024
    NAT dest       4.99.99.100:1024    ->       3.0.0.101:49174
UDP            0.0.0.0:0     ->    3.99.99.100:1024  Forward  I           20999
Gate id: 8590000384
    NAT source         0.0.0.0:0       ->     4.99.99.100:1024
    NAT dest       3.99.99.100:1024    ->       4.0.0.102:49234
UDP          4.0.0.102:0     ->    4.99.99.100:5060  Forward  I               3
Gate id: 4295033089
    NAT source       4.0.0.102:0       ->     3.99.99.100:5060
    NAT dest       4.99.99.100:5060    ->       3.0.0.101:5060
UDP          3.0.0.101:0     ->    3.99.99.100:5060  Forward  I               2
Gate id: 4295033088
    NAT source       3.0.0.101:0       ->     4.99.99.100:5060
    NAT dest       3.99.99.100:5060    ->       4.0.0.102:5060
UDP            0.0.0.0:0     ->    3.99.99.100:1025  Forward  I               0
Gate id: 8590000384
    NAT source         0.0.0.0:0       ->     4.99.99.100:1025
    NAT dest       3.99.99.100:1025    ->       4.0.0.102:49235
UDP          4.0.0.102:0     ->    4.99.99.100:1025  Forward  I               0
Gate id: 8590000385
    NAT source       4.0.0.102:0       ->     3.99.99.100:1025
    NAT dest       4.99.99.100:1025    ->       3.0.0.101:49175
```

### *Displaying Extensive Information About All PGCP Flows*

**Purpose** Display extensive information about all PGCP flows, using the show services pgcp flows extensive command.

**Action**

```
user@host> show services pgcp flows extensive
Interface: sp-1/2/0, Service set: pgcp-svc-set
Flow                                                State    Dir       Frm count
Gate id: 4295033088
UDP                :::0     ->  222::99:99:20:5060  Forward  I              0
    NAT source              :::0       ->     4::99:99:20:5060
    NAT dest     222::99:99:20:5060    ->        4::1:1:3:5060
  Byte count: 0
  Flow role: Master, Timeout: 429496728
  Tman Policing: ON
  SDR    : 10000 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 10000 bytes per second
  PDR MBS: 1000 bytes
Gate id: 4295033088
UDP                :::0     ->    4::99:99:20:5060  Forward  I              0
    NAT source              :::0       ->   222::99:99:20:5060
    NAT dest       4::99:99:20:5060    ->          222::6:5060
  Byte count: 0
  Flow role: Responder, Timeout: 429496728
  Tman Policing: ON
  SDR    : 500 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 500 bytes per second
  PDR MBS: 1000 bytes
```

### *Displaying Extensive Information About PGCP Flows for a Specific Gate*

**Purpose** Display extensive information about PGCP flows for a specific gate using the show services pgcp flows *gateway-name* gate-id *gate-id* extensive command.

**Action**

```
user@host> show services pgcp flows vpg-1 gate-id 4295033088 extensive
Interface: rsp1, Service set: pgcp-svc-set-1
Flow                                                State    Dir       Frm count
Gate id: 4295033088
UDP          4.0.0.102:0     ->    10.50.100.1:1024  Forward  I              0
    NAT source       4.0.0.102:0       ->     20.50.100.1:1024
    NAT dest       10.50.100.1:1024    ->       4.0.0.101:10000
  Byte count: 0
  Flow role: Master, Timeout: 429496728
  Tman Policing: ON
  SDR    : 10000 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 10000 bytes per second
  PDR MBS: 1000 bytes
Gate id: 4295033088
UDP          4.0.0.102:0     ->    10.50.100.1:1025  Forward  I              0
    NAT source       4.0.0.102:0       ->     20.50.100.1:1025
    NAT dest       10.50.100.1:1025    ->       4.0.0.101:10001
  Byte count: 0
  Flow role: Initiator, Timeout: 429496728
```

```
  Tman Policing: ON
  SDR    : 500 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 500 bytes per second
  PDR MBS: 1000 bytes
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Monitoring PGCP Conversations

A PGCP conversation is a group of flows that are grouped based on the call that they belong to. A voice call typically has four flows—an RTP and an RTCP flow in the forward direction, and an RTP and an RTCP flow in the reverse direction. If the PGC does not create RTCP flows, the call has two flows—an RTP flow in each direction.

You can view PGCP conversations with the show services pgcp conversations command. You can view information for all PGCP conversations, or you can include keywords to limit the number of conversations displayed or to display conversations for a particular:

- VPG
- Application protocol
- Source or destination port
- Source or destination prefix
- Protocol
- Service set

### *Displaying All PGCP Conversations*

**Purpose** Display standard information about all PGCP conversations using the show services pgcp conversations command.

**Action**

```
user@host> show services pgcp conversations
Interface: sp-0/3/0, Service set: pgcp-svc-set-1

Conversation: ALG protocol: any
  Number of initiators: 2, Number of responders: 2
Flow                                                State    Dir       Frm count
UDP          4.0.0.102:0     ->    4.99.99.100:1024  Forward  I           20051
Gate id: 8590000385
    NAT source       4.0.0.102:0       ->     3.99.99.100:1024
    NAT dest       4.99.99.100:1024    ->       3.0.0.101:49174
UDP          4.0.0.102:0     ->    4.99.99.100:1025  Forward  I               0
Gate id: 8590000385
    NAT source       4.0.0.102:0       ->     3.99.99.100:1025
    NAT dest       4.99.99.100:1025    ->       3.0.0.101:49175
UDP            0.0.0.0:0     ->    3.99.99.100:1024  Forward  I           19551
Gate id: 8590000384
    NAT source         0.0.0.0:0       ->     4.99.99.100:1024
    NAT dest       3.99.99.100:1024    ->       4.0.0.102:49234
```

```
UDP            0.0.0.0:0     ->    3.99.99.100:1025  Forward  I              0
Gate id: 8590000384
    NAT source         0.0.0.0:0       ->     4.99.99.100:1025
    NAT dest       3.99.99.100:1025    ->       4.0.0.102:49235

Conversation: ALG protocol: any
  Number of initiators: 1, Number of responders: 1
Flow                                                State    Dir       Frm count
UDP          3.0.0.101:0     ->    3.99.99.100:5060  Forward  I              2
Gate id: 4295033088
    NAT source       3.0.0.101:0       ->     4.99.99.100:5060
    NAT dest       3.99.99.100:5060    ->       4.0.0.102:5060
UDP          4.0.0.102:0     ->    4.99.99.100:5060  Forward  I              3
Gate id: 4295033089
    NAT source       4.0.0.102:0       ->     3.99.99.100:5060
    NAT dest       4.99.99.100:5060    ->       3.0.0.101:5060
```

## *Displaying Extensive Information About All PGCP Conversations*

**Purpose** Display extensive information about all PGCP conversations using the show services pgcp conversations extensive command.

**Action**

```
user@host> show services pgcp conversations extensive
Interface: sp-1/2/0, Service set: pgcp-svc-set

Conversation: ALG protocol: any
  Number of initiators: 1, Number of responders: 1
Flow                                                State    Dir       Frm count
Gate id: 4295033088
UDP                :::0     ->  222::99:99:20:5060  Forward  I              0
    NAT source             :::0       ->     4::99:99:20:5060
    NAT dest     222::99:99:20:5060    ->        4::1:1:3:5060
  Byte count: 0
  Flow role: Master, Timeout: 429496728
  Tman Policing: ON
  SDR    : 10000 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 10000 bytes per second
  PDR MBS: 1000 bytes
Gate id: 4295033088
UDP                :::0     ->    4::99:99:20:5060  Forward  I              0
    NAT source             :::0       ->   222::99:99:20:5060
    NAT dest       4::99:99:20:5060    ->          222::6:5060
  Byte count: 0
  Flow role: Responder, Timeout: 429496728
  Tman Policing: ON
  SDR    : 500 bytes per second
  SDR MBS: 1000 bytes
  PDR    : 500 bytes per second
  PDR MBS: 1000 bytes
```

**Related Topics**

- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

Chapter 6

# Managing the Packet Gateway

This chapter describes how to manage the PGCP process and how to shut down VPGs and virtual interfaces. Topics include:



## Managing the PGCP Process

You can stop and start the PGCP process in any or these ways:

- Restart the PGCP process. In this procedure, the PGCP process is considered configured.
- Disable and enable the PGCP process. In this procedure, the PGCP process is considered configured.
- Activate and deactivate the PGCP service.

The packet gateway controller (PGC) cannot send add, modify, or subtract commands to the PGCP process.

### *Restarting the PGCP Process*

**CAUTION:** We recommend that you do not restart the software process unless instructed to do so by a Juniper Networks customer support engineer. A restart might cause the router to drop calls and interrupt transmission.

Three options are available when you restart the PGCP process:

- gracefully—Restart the software process after calls have ended.
- immediately—Immediately restart the software process.

- soft— Reread and reactivate the configuration without completely restarting the software process.

To restart the PGCP process, enter the restart pgcp-service command in operational mode:

```
user@host>restart pgcp-service immediately
```

When you restart the PGCP process, the process is stopped and then restarted.

## *Disabling and Enabling the PGCP Process*

This topic describes the process of disabling and enabling the PGCP process.

### Disabling the PGCP Process

To disable the PGCP process, enter the set system processes pgcp-service disable statement in configuration mode, and then commit your configuration:

```
user@host# set system processes pgcp-service disable
user@host# commit
```

### Enabling the PGCP Process

To enable the PGCP process, enter the delete system processes pgcp-service disable statement in configuration mode, and then commit your configuration:

```
user@host# delete system processes pgcp-service disable
user@host# commit
```

## *Activating and Deactivating PGCP Services*

This topic describes the process of deactivating and activating the PGCP service.

### Deactivating the PGCP Service

To deactivate the PGCP service, enter the deactivate services pgcp statement in configuration mode:

```
user@host# deactivate services pgcp
```

When you enter this statement, the PGCP process is stopped and the PGCP configuration is removed from the router.

### Activating the PGCP Service

The PGCP service is activated by default. If you have deactivated the service, you can activate it by entering the activate services pgcp statement in configuration mode:

```
user@host# activate services pgcp
```

## Shutting Down a VPG

You can shut down the VPG in two ways—forced or graceful:

- Forced—The VPG immediately removes all gates and disconnects from the PGC. The packet gateway does not attempt to establish a new connection.
- Graceful—The VPG goes out of service by entering a draining mode and waiting for all terminations to be subtracted before going out of service. During the draining, the VPG accepts only subtract and audit commands from the PGC.

### *Forcing the Shutdown of a VPG*

To perform a forced shutdown of the VPG, enter set service-state out-of-service-forced at the [edit services pgcp gateway *gateway-name*] hierarchy level. For example:

```
[edit services pgcp gateway vpg-1]
user@host#set service-state out-of-service-forced
```

### *Performing a Graceful Shutdown of a VPG*

To perform a graceful shutdown of the VPG, enter set service-state out-of-service-graceful at the [edit services pgcp gateway *gateway-name*] hierarchy level. For example:

```
[edit services pgcp gateway vpg-1]
user@host#set service-state out-of-service-graceful
```

### *Making the VPG Operational Again*

To cause the VPG to be operational again and available for traffic, set the service state to in-service with the set service-state in-service statement. When the VPG is in service, it attempts to connect to the PGC and accepts all PGCP commands from the PGC. For example:

```
[edit services pgcp gateway vpg-1]
user@host#set service-state in-service
```

## Shutting Down a Virtual Interface

Shutting down the virtual interface is useful when you do not want to shut down the entire VPG. You can shut down the virtual interface two ways—forced or graceful:

- Forced—The virtual interface immediately removes all calls and disconnects from the physical interface.
- Graceful—The virtual interface goes out of service by entering a draining mode and waiting for all terminations to be subtracted before going out of service. During the draining, terminations associated with the virtual interface accept only subtract commands from the PGC.

### *Forcing the Shutdown of a Virtual Interface*

To perform a forced shutdown of a virtual interface, enter set service-state out-of-service-forced at the [edit services pgcp virtual-interface *interface-number*] hierarchy level. For example:

```
[edit services pgcp virtual-interface v1]
user@host#set service-state out-of-service-forced
```

### *Performing a Graceful Shutdown of a Virtual Interface*

To perform a graceful shutdown of a virtual interface, enter set service-state out-of-service-graceful at the [edit services pgcp virtual-interface *interface-number*] hierarchy level. For example:

```
[edit services pgcp virtual-interface v1]
user@host#set service-state out-of-service-graceful
```

### *Making the Virtual Interface Operational Again*

To cause the virtual interface to be operational again and available for traffic, set the service state to in-service with the set service-state in-service statement. When the virtual interface is in service, it is connected to the physical interface and accepts all voice calls. For example:

```
[edit services pgcp virtual-interface v1]
user@host#set service-state in-service
```

## Maintaining Synchronization Between the PG and the PGC

The PGCP software uses the following features to maintain synchronization between the PG and the PGC.

- Detecting Hanging Terminations on page 106
- Detecting PGC Failures on page 108
- Maintaining Synchronization by Auditing Terminations on page 108

### *Detecting Hanging Terminations*

Synchronization of termination information between the PG and the PGC is essential for traffic, maintenance, and charging purposes. If a termination does not exchange messages for a period of time, corresponding data for the termination might be mismatched on the PG and the PGC, and the termination can be hanging. Hanging terminations can consume resources that can be used for chargeable calls.

To detect possible hanging terminations, the PG uses a timer that begins when a message is exchanged for a specific termination. If the termination does not receive a message when the timer expires, the PG notifies the PGC. If the data on the PGC does not match the data on the PG, the PGC returns one of the following error messages to the PG:

```
Error Code #: 411 Name: The transaction refers to an unknown ContextID
Error Code #: 430 Name: Unknown TerminationID
Error Code #: 435 Name: Termination ID is not in specified Context
```

The PGC is responsible for correcting mismatches in data. For example, the PGC can subtract the indicated termination and clear associated contexts. The PGC can also audit the termination service state to check its records before taking further action.

### Activating and Configuring Hanging Termination Detection

The timerx timer sets the interval between the last message exchanged for a specific termination and when the PG sends a notification to the PGC. The timer resets when a message is exchanged on the termination and when the PG notifies the PGC. For example, if the PG notifies the PGC that there has been no activity on the termination, and the PGC does not modify or subtract the indicated termination, the PG again waits the specified time, and sends another notification if there still has not been activity on the termination.

Setting the timer to a value other than zero (0) activates hanging termination detection on new and modified terminations. The timer value that you set is the default value, and it can be overridden by H.248 messages sent from the PGC. To set the timer:

```
[edit services pgcp gateway vpg-1 h248-properties hanging-termination-detection]
user@host#set timerx 30
```

### Deactivating Hanging Termination Detection

To deactivate hanging termination detection, set the timerx statement to 0:

```
[edit services pgcp gateway vpg-1 h248-properties hanging-termination-detection]
user@host#set timerx 0
```

### Displaying the Value of the Timerx Timer Configured on the VPG

**Purpose** Display the value of the timerx that is configured on the VPG

**Action**

```
[edit services pgcp gateway vpg-1 h248-properties hanging-termination-detection]
user@host# show
timerx 30;
```

### Displaying the Value of the Hanging Termination Timer for a Termination

**Purpose** Display the timerx value for a termination.

**Action**

```
user@host> show services pgcp terminations vpg-1 h248
Termination information:
ip/4/vif-0/2 {
      MEDIA {
        TERMINATIONSTATE { SERVICESTATES = INSERVICE },
        STREAM = 1 {
          LOCALCONTROL { MODE = SENDRECEIVE, GM/LSP = 0 },
           LOCAL {
            v = 0,
```

```
            c=IN IP4 10.50.100.1
            m=- 1080 RTP/AVP  –
            b=AS:0
          }
        }
      }
      EVENTS { HANGTERM/THB { TIMERX = 30 } }
}
```

## *Detecting PGC Failures*

The packet gateway supports the inactivity timer package defined in *Gateway control protocol: Inactivity timer package H.248.14, March 2002*. This feature allows the PG to detect the failure of its active PGC through message inactivity. The inactivity timer is applied to the root terminations of a VPG.

The inactivity timer package specifies a default maximum inactivity time. This timer resets each time the PG receives a message from the PGC. If the PG does not receive a message from the PGC before the maximum inactivity time expires, it sends a Notify message with the observed inactivity timeout event to the PGC. If the PGC does not reply to the message, the PG considers the PGC as failed.

### Configuring the Inactivity Timer Package

**Step-by-Step Procedure** To configure the inactivity timer package:

1. Access the inactivity timeout configuration for a VPG.

   ```
   [edit services pgcp]
   user@host#edit gateway vpg-1 h248-properties inactivity-timer inactivity-timeout
   ```

2. Specify whether the PG detects inactivity timeout events received from the PGC by default.

   ```
   [edit services pgcp gateway vpg-1 h248-properties inactivity-timer inactivity-timeout]
   user@host#set detect
   ```

3. Specify a default value for the maximum inactivity time. The default value is used if the PGC requests that the PG detect the inactivity timeout event, but the PGC does not set a value for the maximum inactivity time.

   ```
   [edit services pgcp gateway vpg-1 h248-properties inactivity-timer inactivity-timeout]
   user@host#set maximum-inactivity-time default 24000
   ```

## *Maintaining Synchronization by Auditing Terminations*

You can use the AuditValue command in H.248 messages to maintain synchronization between the PG and the PGC. The AuditValue command requests the current values of a descriptor or of a single property, event, signal, or statistic associated with terminations and contexts. You can include selection criteria in the AuditValue command to filter the returned values.

The software supports the following characters for audit selection criteria:

- equal to (=)
- not equal to (#)
- less than (<)
- greater than (>)

### Using AND/OR Logic with Audit Commands

You can include multiple criteria with an AND or an OR logic operation (ANDLgc, ORLgc) to indicate how the selection criteria are interpreted. If you do not include a logic operation, AND logic operation is applied.

### Example: Audit Section Filter with AND Logic

**Purpose** Create an audit selection filter for contexts and terminations that have the saf property set to on AND the spr property set to less than 4075.

**Action**

```
Transaction = 201 {
    Context = * {
        ContextAudit {
            ANDLgc
        },
        AuditValue = * {
            Audit {
                Media {
                    LocalControl {
                        gm/saf=on,
                        gm/spr<4075,
                    }
                }
            }
        }
    }
}
```

**Meaning** The result of this audit is a list of terminations that have the saf property set to on and the spr property set to less than 4075.

### Example: Audit Section Filter with OR Logic

**Purpose** Create an audit selection filter for contexts and terminations that have the spf property set to on OR the sam property set to 10.10.0.3.

**Action**

```
Transaction = 202 {
    Context = * {
        ContextAudit {
            ORLgc
        },
        AuditValue = * {
            Audit {
                Media {
```

```
                    LocalControl {
                        gm/spf=on,
                        gm/sam=10.10.0.3,
                    }
                }
            }
        }
    }
}
```

**Meaning** The result of this audit is a list of terminations that have the spf property set to on or the sam property set to 10.10.0.3.

## Managing Overload Control for New Call Setup

You can enable the VPG to notify a PGC when it experiences processing overload that could prevent the timely execution of H.248 transactions. The PGC can then adjust the rate at which H.248 transactions are passed to the VPG.

The PGC activates the overload control feature by setting the overload control event. When overload control is active, incoming H.248 transactions are pushed into a work queue. The VPG sends an overload notification to the PGC when the content of the work queue reaches a configured limit, defined as a percentage of the maximum queue size (queue-limit-percentage). The VPG then sends an overload notification for each received ADD command in incoming H.248 transactions. The PGC should lower the rate used to admit calls to the VPG (the admitted rate). This lower rate is configured on the PGC. When transactions in the work queue occupy 100 percent of the work queue's maximum size, the VPG drops received transactions and sends error code 510 (Insufficient resources). When transactions in the work queue occupy less than the queue-limit-percentage, overload notifications are no longer sent.

### *Configuring Overload Control for Voice Calls*

You configure overload control by configuring the queue limit percentage.

**Step-by-Step Procedure** To configure overload-control for voice calls:

1. Access the configuration of overload control.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#edit overload-control
   ```

2. Configure a queue limit percentage.

   ```
   [edit services pgcp gateway vpg-1 overload-control]
   user@host#set queue-limit-percentage 85
   ```

**NOTE:** Only one work queue exists for the entire device. The queue limit percentage that you configure applies to the entire device, not just the VPG named in the hierarchy level you have chosen. If you enter multiple queue limit percentages for different VPGs, only the last entry is used.

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Preventing Excessive Media Inactivity Notifications

You can prevent a PG from sending an excessive number, or *avalanche*, of media inactivity notifications to the PGC in a short period of time. Such an avalanche can severely degrade the performance of the PGC. By default, all media inactivity notifications are sent immediately from the PG to the PGC. When an upstream device in the network fails, all existing terminations in the PG report media inactivity at about the same time. You can configure H.248 notification behavior to regulate the flow of notifications sent to the PGC.

Regulated notification is activated either by a request from the PGC or by CLI commands. When you choose regulated notification you can explicitly configure the frequency for notifications of media inactivity events as follows:

- Send only one media inactivity notification. This option is not available as part of the H.248 notification behavior package.
- Send a configurable percentage of media inactivity notifications ( 0 through 100).

**Notification History—**The Notification Behavior Package only limits the sending of notifications to the PGC, not the occurrences of the event. We recommend you enable the recording of all media inactivity notifications for future retrieval and, optionally, request that the inactivity notifications time stamps are also enabled for future retrieval.

### *Configuring H.248 Notification Behavior to Prevent Excessive Media Inactivity Notifications*

By properly setting H.248 notification behavior properties, you can prevent a PG from sending an avalanche of media inactivity notifications that can flood the PGC and adversely affect the processing of H.248 transactions. You can also enable recording of all media inactivity events for future retrieval.

The H.248 Notification Behavior package is defined in annex E.15 of the *Gateway control protocol v.3, ITU-T Recommendation H.248.1, September, 2005*.

**Step-by-Step Procedure** To configure default values for H.248 Notification Behavior properties:

1. Access the configuration of H.248 application-data-inactivity-detection.

   ```
   [edit services pgcp gateway vpg-1]
   user@host#edit h248-properties application-data-inactivity-detection
   ```

2. Turn notification on.

   ```
   [edit services pgcp gateway vpg-1 h248-properties
     application-data-inactivity-detection]
   user@host#set ip-flow-stop-detection regulated-notify
   ```

3. Access the configuration of H.248 notification behavior.

```
[edit services pgcp gateway vpg-1]
user@host#edit h248-properties notification-behavior
```

4. Configure the default frequency of notification messages for the media inactivity event. The PGC can override this default by requesting a different frequency. If you specify **once**, only one notification is sent and the PGC cannot retrieve the notification by use of an audit for the root.termination.

```
[edit services pgcp gateway vpg-1 h248-properties notification-behavior]
user@host#set notification-regulation default 10
```

5. Enable retrieval of event notifications by the PGC at the [edit services pgcp gateway *gateway-name* h248-options] level.

```
[edit services pgcp gateway vpg-1 h248-properties notification-behavior]
user@host#edit services pgcp gateway vpg-1 h248–options
[edit services pgcp gateway vpg-1 h248-options]
user@host#set audit-observed-events-returns-history;
```

6. Enable retrieval of time stamps of recorded notifications by the PGC.

```
[edit services pgcp gateway vpg-1 h248-options]
user@host#edit services pgcp gateway vpg-1 h248-properties
[edit services pgcp gateway vpg-1 h248-properties]
user@host#set event-timestamp-notification request-timestamp requested
```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## *Managing the Rate for All Notifications Sent by a PIC*

You can limit the rate at which the PIC send all notifications. Using this functionality, you can prevent an avalanche of media inactivity notifications and throttle all notifications to a configured rate. Doing this enables the system to maintain a stable state when the PIC is generating a large volume of messages.

### Limiting The Rate for All Notifications from a PIC

You can configure the aggregate rate of notifications coming from a PIC. Using this approach, you limit the rate for all messages from the PIC, not just notifications.

**Step-by-Step Procedure**

To configure a rate limit for a PIC to send notifications to the PGC:

1. Access the configuration of PGCP parameters.

```
[edit services]
user@host#edit pgcp
```

2. Configure the rate (per second) for PICs to send messages to the PGC.

```
[edit services pgcp ]
user@host#set notification-rate-limit 25
```

**Related Topics**
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Controlling ServiceChange Commands Sent from the VPG to the PGC

For seamless interoperability between the JUNOS packet gateway and PGC devices, you can control the method and reason that the VPG includes in ServiceChange commands that it sends to the PGC. You can also prevent the system from being overloaded with messages for certain state changes by specifying that the VPG not send a request or notification when those changes occur.

You can specify the method and reason that the VPG includes in ServiceChange commands when the state of one of the following changes:

- Control associations
- Virtual interfaces
- Contexts

### *Control Association States*

A control association is a relationship where the PGC is controlling the VPG. Each VPG has only one control association at any time.

Table 10 on page 113 describes the control association states.

**Table 10: Control Association States**

| Control Association State | Description |
|---|---|
| Disconnected | The control association is in the Disconnected state. No PGC is controlling the VPG, and incoming H.248 messages are ignored. The control association remains disconnected as long as the VPG is Out-of-Service.<br><br>Depending on what caused the VPG to become Out-of-Service, the VPG either drops H.248 commands or answers them with a port-unreachable ICMP error. |
| Connecting | The control association is in Connecting state between the time the VPG sends a registration request to the PGC and the time the PGC accepts, rejects, or aborts the request.<br><br>The VPG rejects incoming H.248 commands while the control association is in the Connecting state with error # 505: “Transaction Request Received before a ServiceChange Reply has been received”. |
| Draining | The control association enters the Draining state when an administrator instructs the VPG to gracefully transition from In-Service to Out-of-Service. The PGC transitions to Out-of-Service when the controlling PGC subtracts all of the VPG’s H.248 terminations.<br><br>The VPG accepts only Subtract and AuditValue commands from the controlling PGC. It rejects all other commands with error # 502: “Not Ready”. |

When the state of a control association changes, the VPG can send the following types of ServiceChange commands to the PGC:

- Registration Requests—The VPG sends a Registration Request ServiceChange command to request that a PGC become its controlling PGC. The VPG sends these requests when a control association enters the Connecting state.
- Unregistration Messages—The VPG sends an Unregistration ServiceChange command to its controlling PGC when it transitions to the Out-of-Service (Disconnected) service state because of an administration operation or a failure. The failure can be the result of a MultiServices PIC or Flexible PIC Concentrator (FPC) failure or because the MultiServices PIC was powered off or removed.
- Notification Messages—The VPG sends a notification ServiceChange command to its controlling PGC when the control association transitions between the Connected and Draining states and vice versa.

### Method and Reason Options for Control Association State Changes

You can control the ServiceStateMethod and ServiceStateReason that the VPG includes in ServiceChange commands for control associations.

You can use the CLI to specify the method and reason that the VPG includes in ServiceChange commands for control associations. Table 11 on page 114 shows the method and reason options available for each reported state and the events that led to the report.

**Table 11: Options for Method and Reason in ServiceChange Commands for Control Associations**

| Reported Association State | Event Leading to Report | Options | Embedded H.248 Reason | Explanation |
|---|---|---|---|---|
| Disconnect | Controller failure | FL/909 | PGC impending failure | VPG is reregistering with a new PGC following a disconnection of the VPG and PGC. |
| | | RS/902 | Warm boot | VPG is reregistering with a new PGC following a disconnection of the VPG and PGC. |
| | Reconnect | DC/900 | Service restored | VPG is registering with the last controlling PGC following a disconnection of the VPG and PGC. |
| | | RS/902 | Warm boot | VPG is transitioning to In-Service, and the previously installed state is retained. |

**Table 11: Options for Method and Reason in ServiceChange Commands for Control Associations** *(continued)*

| Reported Association State | Event Leading to Report | Options | Embedded H.248 Reason | Explanation |
|---|---|---|---|---|
| Down | Administrative | FO/905 | Termination taken out of service | VPG is transitioning to Out-of-Service because of an administrative operation. |
| | | FO/908 | VPG impending failure | VPG root termination transitioned to Out-of-Service and is unable to process request. |
| | | none | | No message is sent for this event. |
| | Failure | FO/904 | Termination malfunctioning | VPG is transitioning to Out-of-Service because of a failure. |
| | | FO/908 | VPG impending failure | VPG root termination transitioned to Out-of-Service because of a failure. |
| | | none | | No message is sent for this event. |
| | Graceful | GR/905 | Termination taken out of service | The control association entered the Draining state because of an administrative operation. |
| | | none | | No message is sent for this event. |
| Up | Cancel graceful | RS/908 | Cancel graceful | The control association transitioned from the Draining state to the Forwarding state. |
| | | none | | No message is sent for this event. |
| | Cold failover | FL/920 | Cold failover | VPG is registering following a graceful Routing Engine switchover. The previously installed state is reset. |
| | | RS/901 | Cold boot | VPG is transitioning to In-Service. The previously installed state is not retained. |
| | Warm failover | FL/919 | PGC impending failure | VPG is registering with a new PGC following a disconnection of the VPG and PGC. |
| | | RS/902 | Warm boot | VPG is transitioning to In-Service, and the previously installed state is retained. |

## Configuring the Method and Reason in ServiceChange Commands for Control Associations

**Step-by-Step Procedure**

To configure the method and reason in ServiceChange commands for control associations:

1. Access the configuration of the service change control association indications properties.

```
[edit services pgcp]
user@host#edit gateway vpg-1 h248-options service-change
  control-association-indications
```

2. Specify the method and reason that the VPG includes in Registration Request ServiceChange commands when it attempts to reregister with the PGC or register with a new PGC after the control association is disconnected.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set disconnect controller-failure restart-902
```

3. Specify the method and reason that the VPG includes in Registration Request ServiceChange commands when it attempts to reregister with the PGC or register with a new PGC after the control association is disconnected.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set disconnect reconnect restart-902
```

4. Specify the method and reason that the VPG includes in Unregistration Messages in ServiceChange commands that it sends to the PGC when a control association transitions to Out-of-Service because of an administrative operation.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set down administrative forced-908
```

5. Specify the method and reason that the VPG includes in Unregistration Messages in ServiceChange commands that it sends to the PGC when a control association transitions to Out-of-Service because of a failure.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set down failure forced-904
```

6. Specify the method and reason that the VPG includes in Notification ServiceChange commands that it sends to the PGC when the control association transitions from In-Service to Out-of-Service-Graceful.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set down graceful graceful-905
```

7. Specify the method and reason that the VPG includes in Notification ServiceChange commands that it sends to the PGC when the control association has returned to the Connected state.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  control-association-indications]
user@host#set up cancel-graceful restart-918
```

8. Specify the method and reason that the VPG includes in Registration ServiceChange commands when it attempts to register with a new PGC following a cold failover.

```
[edit services pgcp gateway vpg-1 h248-options service-change
   control-association-indications]
user@host#set up failover-cold failover-920
```

9. Specify the method and reason that the VPG includes in Registration ServiceChange commands when it attempts to register with a new PGC following a warm failover.

```
[edit services pgcp gateway vpg-1 h248-options service-change
   control-association-indications]
user@host#set up failover-warm restart-902
```

## *Virtual Interface States*

Table 12 on page 117 describes the virtual interface states.

**Table 12: Virtual Interface States**

| Virtual Interface Operational State | Description |
|---|---|
| Blocked | A virtual interface is in the Blocked state when the interface is Out-of-Service.<br><br>While a virtual interface is in the Blocked state, all VPGs do not add new terminations using the interface. Likewise, the VPGs rejects H.248 commands other than Subtract and AuditValue commands on existing terminations on the interface. The error that the VPG returns for commands that it rejects depends on the reason that caused the virtual interface to be Out-of-Service:<br><br>■ Error #502: "Not ready"—If the virtual interface is Out-of-Service because of an administrative operation.<br>■ Error #529: "Internal hardware failure in PG"—If the virtual interface is Out-of-Service because of a failure. |
| Forwarding | A virtual interface is in the Forwarding state when it is functioning normally. All gates are using the interface process data flows according to the H.248 properties installed on them. |
| Draining | A virtual interface enters the Draining state when an administrator instructs the virtual interface to gracefully transition from In-Service to Out-of-Service. The virtual interface automatically transitions to Out-of-Service when it is no longer used by any termination in any of the VPGs.<br><br>A virtual interface that is in the Draining state is In-Service and existing gates process data flows normally. However, as in the Blocked state, the VPGs do not add new terminations using that virtual interface or to perform any command other than Subtract and AuditValue on existing terminations on the interface. If the VPG receives other commands, it replies with error #502: "Not ready". |

When the state of a virtual interface changes, the VPG can send the following types of ServiceChange commands to the PGC:

■ Service-Restoration—The VPG sends Service-Restoration ServiceChange commands when a virtual interface transitions to the In-Service service state; that is, it transitions from the Blocked state to the Forwarding operational state.

- Service-Interruption—The VPG sends Service-Interruption ServiceChange commands when a virtual interface transitions to the Out-of-Service, or Blocked, service state.
- Notification Messages—The VPG sends Notification ServiceChange commands when a virtual interface transitions between the Forwarding and Draining states and vice versa.

### Method and Reason Options for Virtual Interface State Changes

You can control the ServiceStateMethod and ServiceStateReason that the VPG includes in ServiceChange commands for virtual interface state changes.

You can use the CLI to specify the method and reason that the VPG includes in ServiceChange commands for virtual interfaces. Table 13 on page 118 explains the method and reason options available for each reported state and the events that led to the report.

**Table 13: Options for Method and Reason in ServiceChange Commands for Virtual Interfaces**

| Reported State | Event Leading to Report | Options | Embedded H.248 Reason | Explanation |
|---|---|---|---|---|
| Virtual interface down | Administrative | FO/905 | Termination taken out of service | Virtual interface is transitioning to Out-of-Service because of an administrative operation. |
| | | FO/906 | Loss of lower-layer connectivity | Virtual interface is transitioning to Out-of-Service because of a loss of layer 2 connectivity caused by the logical or physical interface being administratively disabled. |
| | | none | | No message is sent for this event. |
| | Failure | FO/904 | Termination malfunctioning | Virtual interface is transitioning to Out-of-Service because of an internal failure. |
| | | FO/906 | Loss of lower-layer connectivity | Virtual interface is transitioning to Out-of-Service because of a loss of layer 2 connectivity on the logical or physical interface. |
| | | none | | No message is sent for this event. |
| | Graceful | GR/905 | Termination taken out of service | Virtual interface has entered the Draining state. |
| | | none | | No message is sent for this event. |
| | Link loss | FO/906 | Loss of lower-layer connectivity | Virtual interface is transitioning to Out-of-Service because of a loss of layer 2 connectivity on the logical or physical interface. |
| | | none | | No message is sent for this event. |

**Table 13: Options for Method and Reason in ServiceChange Commands for Virtual Interfaces** *(continued)*

| Reported State | Event Leading to Report | Options | Embedded H.248 Reason | Explanation |
|---|---|---|---|---|
| Virtual interface up | Cancel graceful | RS/918 | Cancel graceful | Virtual interface has returned to the Forwarding state. |
| | | none | | No message is sent for this event. |
| | warm | RS/900 | Service restored | Virtual interface has become In-Service and is in the Forwarding state. |
| | | none | | No message is sent for this event. |

## Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces

**Step-by-Step Procedure**

To configure the method and reason in ServiceChange commands for virtual interfaces:

1. Access the configuration of the service change virtual interface indications properties.

```
[edit services pgcp]
user@host#edit gateway vpg-1 h248-options service-change
   virtual-interface-indications
```

2. Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when a virtual interface changes to Out-of-Service because of an administrative operation.

```
[edit services pgcp gateway vpg-1 h248-options service-change
   virtual-interface-indications]
user@host#set virtual-interface-down administrative forced-906
```

3. Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when a virtual interface transitions to Out-of-Service because of a failure.

```
[edit services pgcp gateway vpg-1 h248-options service-change
   virtual-interface-indications]
user@host#set virtual-interface-down failure forced-906
```

4. Specify the method and reason that the VPG includes in Notification ServiceChange commands that it sends to the PGC when a virtual interface transitions between the Forwarding and Draining states.

```
[edit services pgcp gateway vpg-1 h248-options service-change
   virtual-interface-indications]
user@host#set virtual-interface-down graceful none
```

5. Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when the virtual interface transitions to Out-of-Service because of a link loss.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  virtual-interface-indications]
user@host#set virtual-interface-down link-loss forced-906
```

6. Specify the method and reason that the VPG includes in Notification ServiceChange commands that it sends to the PGC when the virtual interface transitions from the Draining state to the Forwarding state.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  virtual-interface-indications]
user@host#set virtual-interface-up cancel-graceful restart-918
```

7. Specify the method and reason that the VPG includes in Service-Restoration ServiceChange commands that it sends to the PGC when a virtual interface transitions to In-Service.

```
[edit services pgcp gateway vpg-1 h248-options service-change
  virtual-interface-indications]
user@host#set virtual-interface-up warm restart-900
```

## *Context States*

The VPG sends context Service-Interruption messages when the gates of a specific context no longer provide their configured service. When such a message is issued, both terminations included in the context become Out-of-Service.

You can use the CLI to specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when a state loss occurs.Table 14 on page 120 describes the method and reason options available.

**Table 14: Options for Method and Reason in ServiceChange Commands for Specific Contexts**

| Reported State | Event Leading to Report | Options | Embedded H.248 Reason | Explanation |
|---|---|---|---|---|
| State loss | Mismatch between pgcpd process and MultiServices PIC states | FO/910 | State loss because of a media failure | A mismatch between the pgcpd process and the MultiServices PIC states on one or more of the context's gates. |
| | | FO/915 | State loss | A mismatch between the pgcpd process and the MultiServices PIC states was detected on one or more of the context's gates. |
| | | none | | No message is sent for this event. |

### Configuring the Method and Reason in ServiceChange Commands for Contexts

**Step-by-Step Procedure** To configure the method and reason in ServiceChange commands for contexts:

1. Access the configuration of the service change context indications properties.

```
[edit services pgcp]
```

```
user@host#edit gateway vpg-1 h248-options service-change context-indications
```

2. Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC after a state loss on a specific context.

```
[edit services pgcp gateway vpg-1 h248-options service-change context-indications]
user@host#set state-loss forced-915
```

Chapter 7

# Maintenance and Failover in the Packet Gateway

Topics include:



## Maintenance and Failover in the Packet Gateway Overview

By providing redundancy for both the Routing Engine and the service PICs, the packet gateway high-availability (HA) architecture allows for a multilevel redundant solution. This solution ensures service and call continuity, which is necessity for VoIP services.

Figure 24 on page 123 shows the packet gateway HA architecture.

**Figure 24: Packet Gateway HA Architecture**

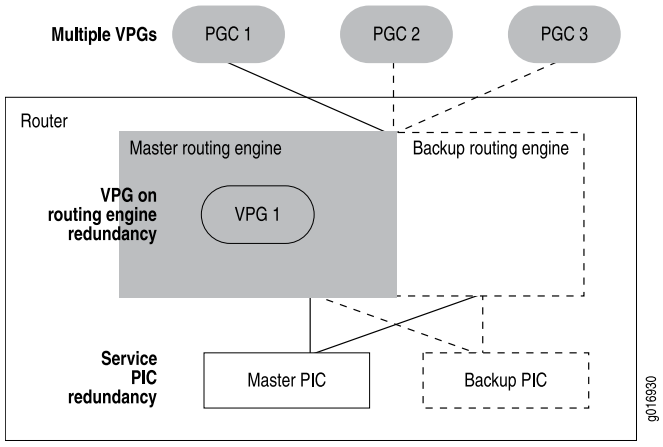


Table 15 on page 124 indicates how service PIC failure and Routing Engine failure affect different types of services.

**Table 15: How Service PIC and Router Engine Failures Affect Service and Call Continuity**

| Service | Service PIC Failure | Routing Engine Failure |
|---|---|---|
| Stable H.248 contexts | Call continuity—All calls are lost.<br><br>Service continuity—Provided when the backup service PIC is available. | Call continuity—All existing calls are maintained through the failure<br><br>Service continuity—New calls and session can be established when the backup Routing Engine takes control. |
| H.248 contexts are modified | Call continuity—All existing call states are lost.<br><br>Service continuity—Calls and sessions that are in a Setup state (gates and terminations not fully defined) are lost | Call continuity—All existing call states are lost.<br><br>Service continuity—Calls and sessions that are in a Setup state are lost. |
| Connection to the PGC | Call continuity—All calls are lost.<br><br>Service continuity—Provided when the backup service PIC is available. | Call continuity—No calls are lost.<br><br>Service continuity—Provided when the backup Routing Engine takes control. |

## Failover in Case of a Routing Engine Failure

Graceful Routing Engine switchover (GRES) is supported in case of a Routing Engine failure. A failure of the Routing Engine stops the PGCP process. When the JUNOS software high-availability framework detects the Routing Engine failure, it switches control of the packet gateway to the PGCP process on the backup Routing Engine. The PGCP process stores completed H.248 states, and these states survive a restart or a switchover. After the restart or switchover is completed, the PGCP process reapplies the states.

### *Gate Synchronization Procedure*

When the IPC connection between the PGCP process and the MultiServices PIC is being reestablished, a synchronization process restores the previous state of gates, so session and call continuity is achieved. That is, existing calls stay active and voice sessions are not disconnected.

The synchronization process can result in the mismatching of gates. Mismatched gates are handled as follows:

- If the PGCP process detects gates that exist in the Routing Engine, but are missing in the PIC, the PGCP process reinstalls the gates on the PIC. Existing calls receive service from the PIC when the synchronization is complete.
- If the PGCP process detects gates that exist on the PIC, but not on the Routing Engine, the PGCP process removes the gates from the PIC. This removal causes the associated sessions to be closed, and the session resources to be released.
- If the PGCP process detects gates that exist on the PIC and on the Routing Engine, but the versions of the gates do not match, the PGCP process forces the Routing Engine version. Doing so maintains call continuity by making sure that the databases are synchronized, while limiting the effect to existing calls and sessions.

You can control the number of mismatches by configuring the synchronization properties.

### *Configuring Synchronization Properties*

You can configure the maximum number of mismatches allowed during the synchronization process between the PIC and the PGCP process.

**Step-by-Step Procedure** To configure properties for synchronization between the PIC and the PGCP process:

1. Access the graceful restart configuration.

   ```
   [edit services pgcp]
   edit gateway vpg-1 graceful-restart
   ```

2. Configure the maximum number of mismatches allowed during the synchronization procedure. If the number of mismatches exceeds this number, the PGCP process clears the state of the PIC and the state of the PGCP process. All calls and sessions are terminated and existing resources are released.

   ```
   [edit services pgcp gateway vpg-1 graceful-restart]
   user@host#set maximum-synchronization-mismatches 20
   ```

3. Configure the number of milliseconds within which you want the synchronization process to complete. If the process is not complete when this time expires, the PGCP process clears the state of the PIC and the state of the PGCP process.

   ```
   [edit services pgcp gateway vpg-1 graceful-restart]
   user@host#set maximum-synchronization-time 300
   ```

### *Displaying the Status of the Routing Engine Synchronization*

**Purpose** To determine the status of the synchronization between the main Routing Engine and the backup Routing Engine, display the value of the replication socket field.

**Action**

```
user@host> show services pgcp active-configuration
. . .
Packet gateway configuration:
  Name                                  :  vpg-1
  IP address                            :  10.10.10.1
  Port                                  :  2944
  Status                                :  In-Service (Registered)
  Active gateway controller             :  dt3
  Cleanup timeout [secs]                :  3600
  Gate inactivity delay [secs]          :  240
  Gate inactivity duration (Q-MI ) [secs]  :  86400
  Replication socket                    :  Disconnected

. . .
```

**Meaning** The replication socket can be in one of the following states:

- Connected (Ready)—The replication is ready, and it is safe to perform a switchover.
- Connected (Syncing)—The replication is synchronizing. It is not safe to perform a switchover.
- Connected (Error)—An error occurred in the previous switchover.
- Disconnected—The backup Routing Engine is down. It is not safe to perform a switchover because there is no backup Routing Engine.

**Related Topics**

- *Part 3, Graceful Routing Engine Switchover* in *JUNOS High Availability Configuration Guide*
- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

## Failover of the Service PICs

The PIC failover procedure for the packet gateway assures service continuity in case of a service PIC failure. The packet gateway architecture provides both a 1:1 redundancy model and a 1:N redundancy model for service PICs.

In the 1:1 redundancy model, there is one primary service PIC and one secondary service PIC that acts as a backup. If the primary service PIC fails, the Routing Engine allocates the secondary service PIC, and the software switches over to the secondary service PIC. The PGCP process continues to provide the PGCP services as if the service PIC has restarted. All states and sessions are lost, but new calls are accepted. When the failed service PIC recovers, it does not take over from the redundant PIC.

In the 1:N redundancy model, you can define one PIC as the secondary of many primary PICs. In this model, after the secondary PIC becomes active, all other primary PICs are left without a secondary PIC. Even when the primary PIC recovers, it does not become a redundant PIC to all of the primary PICs. A recovered primary PIC can replace only the same secondary PIC that previously replaced it. This functionality means that administrative involvement is usually required after a failover event happens.

### *Procedure in Case of Service PIC Failure*

If a service PIC fails, the following procedure takes place:

1. Active calls are lost.
2. The PG notifies the PGC of the failure using an FO/904 ServiceChange message.
3. The PG receives an acknowledgment for the FO/904 message from the PGC.
4. The redundant service PIC (rsp) mechanism allocates the secondary PIC and makes the secondary PIC the new primary PIC.

5. The new primary PIC establishes the IPC connection to the PGCP process on the routing engine.
6. The PGCP process issues an RS/902 registration message to the PGC.
7. The PGCP process receives an acknowledgment of the registration message from the PGC.
8. The PG is ready to accept and process new H.248 commands.

## *Configuring the Packet Gateway for PIC Redundancy*

To configure PIC failover, you configure a redundancy services PIC (rsp) interface that specifies which service PIC is the primary PIC and which service PIC is the secondary PIC. In the service set configuration for the PGCP service, the service set points to the rsp interface as the next-hop service interface.

### Configuring the Redundancy Services PIC (rsp) Interface

To configure, create a redundancy services PIC (rsp) interface, and specify the primary and secondary service PIC and the inside and outside service domains.

**Step-by-Step Procedure** To configure the rsp interface:

1. Configure the interface, and enter edit mode for the interface.

   ```
   [edit]
   user@host#edit interfaces rsp1
   ```

2. Specify the service PIC that is to be the primary PIC.

   ```
   [edit interfaces rsp1]
   user@host#set redundancy-options primary sp-1/2/0
   ```

3. Specify the service PIC that is to be the secondary PIC.

   ```
   [edit interfaces rsp1]
   user@host#set redundancy-options secondary sp-1/3/0
   ```

4. Configure a logical unit and specify the protocol family.

   ```
   [edit interfaces rsp1]
   user@host#set unit 10 family inet
   ```

5. Set the service domain of the logical unit to inside. This unit number must match the unit number of the inside service interface configured in the service set.

   ```
   [edit interfaces rsp1]
   user@host#set unit 10 service-domain inside
   ```

6. Configure another logical unit and specify the protocol family.

   ```
   [edit interfaces rsp1]
   user@host#set unit 20 family inet
   ```

7. Set the service domain of the logical unit to outside. This unit number must match the unit number of the outside service interface configured in the service set.

```
[edit interfaces rsp1]
user@host#set unit 20 service-domain outside
```

### Configuring the Service Set for Redundant Service PICS

When you configure your service set, specify the rsp interface as the next-hop service.

**Step-by-Step Procedure** To configure a service set for redundancy:

1. Create a service set configuration.

```
[edit services]
user@host#edit service-set pgcp-svc-set
```

2. Configure service set as a next-hop service set.

```
[edit services service-set pgcp-svc-set]
user@host#edit next-hop-service
```

3. Specify the rsp interface to the inside network. This unit number must match the unit number of the inside service domain configured in the rsp interface.

```
[edit services service-set pgcp-svc-set next-hop-service]
user@host#set inside-service-interface rsp1.10
```

4. Specify the rsp interface to the outside network. This unit number must match the unit number of the outside service domain configured in the rsp interface.

```
[edit services service-set pgcp-svc-set next-hop-service]
user@host#set outside-service-interface rsp1.20
```

## *Manually Switching from the Primary PIC to the Secondary PIC*

**Purpose** Manually switch from the primary PIC to the secondary PIC.

**Action**

```
user@host>  request interface rsp1 switchover
request succeeded
```

## *Manually Reverting from the Secondary PIC to the Primary PIC*

**Purpose** Manually revert from the secondary PIC to the primary PIC.

**Action**

```
user@host> request interface rsp1 revert
request succeeded
```

## *Displaying the Status of the Redundant Service PICs*

**Purpose** Display the status of the redundant service PICs.

**Action**

```
user@host> show interface redundancy rsp1
Interface  State           Last change  Primary    Secondary  Current status
rsp1       On primary     00:25:18        sp-1/2/0    sp-1/3/0        both up
```

**Related Topics**

- *Chapter 24, Interface Configuration Guidelines* in *JUNOS Services Interfaces Configuration Guide*
- *Chapter 19, Adaptive Services Interface Operational Mode Commands* in *JUNOS Interfaces Command Reference*

Chapter 8

# Troubleshooting the Voice Solution

This chapter explains how to set up trace options and the logging of H.248 messages. Topics include:



## Tracing PGCP Operations

You can trace the following PGCP operations and record them in a log file:

- all—All PGCP operations
- configuration—Configuration events
- debug—Debug messages
- error—Error messages
- gate—Gate request and reply events
- media-function—Media function events
- pgc-connection—PGC connection events
- pgcp-stack-debug—PGCP stack debug events
- pgcp-stack-h248—H.248 messages between VPGs and PGCs
- pgcp-stack-trace—PGCP stack function call events
- routing-socket—Routing-socket events
- trace—Process function calls

All log files are placed in the /var/log directory. When a trace file reaches its maximum size, a .0 is appended to the file name, then a new file is created with a .1 appended, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

**Step-by-Step Procedure** To configure tracing of PGCP operations:

1. Access the PGCP traceoptions configuration.

```
[edit]
user@host# edit services pgcp traceoptions
[edit services pgcp traceoptions]
```

2. Specify a name for the trace file.

```
[edit services pgcp traceoptions]
user@host# set file pgcp
```

3. Set the maximum number of trace files. If you specify a maximum number of files, you also must specify a maximum file size.

```
[edit services pgcp traceoptions]
user@host# set files 10
```

4. Specify the operations that you want to include in the log file. For example:

```
[edit services pgcp traceoptions]
user@host# set flag gate
user@host# set flag pgc-connection
user@host# set flag pgcp-stack-h248
```

5. Prevent any user from reading the log file.

```
[edit services pgcp traceoptions]
user@host# set no-world-readable
```

6. Allow any user to read the log file.

```
[edit services pgcp traceoptions]
user@host# set world-readable
```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

## Logging H.248 Messages

You can use the traceoptions feature to log H.248 messages exchanged between a VPG and its associated PGC. These messages include all transactions, transaction replies, and transaction acknowledgements. Corrupted messages, messages with invalid syntax, and messages that fail interim AH validation are also logged. The captured data includes both the raw H.248 message (as it appears on the wire) and associated metadata.

The captured data is stored by default in the pgcpd file in the /var/log directory. In addition to H.248 messages, this file contains all other PGCP operations that you have chosen to trace. To make it easier to find H.248 messages in the file, there is a BEGIN field and an END field, each with a matching sequence number.

When you enable logging of H.248 messages, the software logs messages for all VPGs on the router. To determine which VPG exchanged the message, use the 5–tuple field in the H.248 message.

### *Fields in the H.248 Messages*

Each H.248 message contains the following fields:

```
TRACE_TIMESTAMP
[BEGIN #SEQNUM  TIMESTAMP  5-TUPLE  DIRECTION  MSG_SIZE]
MegacoMessage
[END #SEQNUM]
```

The following is a sample ServiceChange message:

```
Oct 10 09:32:28
[BEGIN #778  Oct 10 09:32:28  UDP:10.50.40.100:2944->172.16.1.1:2944  TX  168]
AU=0x00000000:0x00000000:0x00000000000000000000000000
!/1 [10.50.40.100]:2944
T=1{
C=-{
SC=ROOT{
SV{MT=RS,RE="901",AD=2944,PF=JNPR_PGCP/1,V=3,20071010T09322800}}}}
[END #778]
```

Table 16 on page 133 describes the fields in H.248 messages.

**Table 16: Description of Fields in H.248 Messages**

| Field | Description |
|---|---|
| TRACE_TIMESTAMP | Time the messaged was captured. The timestamp is in the format *HH:MM:SS.mmmmmm* |
| BEGIN | Constant string BEGIN that indicates the beginning of a message. |
| #SEQNUM | Sequential number of the message used to correlate the BEGIN header with the END footer. |
| TIMESTAMP | Date and time of the message. Time is displayed in the format *HH:MM:SS.mmmmmm* |
| 5-TUPLE | UDP or TCP packet's 5-tuple (transport protocol, source address and port, destination address and port). Used to identify the VPG and the PGC that exchanged the message. |
| DIRECTION | Direction message was sent from the VPG:<br>■ TX for an outgoing H.248 message.<br>■ RX for an incoming H.248 message. |
| MSG_SIZE | Size of message in bytes excluding transport protocol (UDP/TCP/IP) overhead. |
| MegacoMessage | Captured data including the interim authentication header (if any). |
| END | Constant string END that indicates the end of a captured message. |

### *Messages That Exceed Output Buffer Limit*

The trace file has a limit on the size of a single output buffer. Messages longer than this limit are divided into several trace buffers. Messages divided into buffers use the following template:

```
TRACE_TIMESTAMP
  [BEGIN #SEQNUM  TIMESTAMP  5-TUPLE  DIRECTION  MSG_SIZE]
  MegacoMessage (1st part)
  TRACE_TIMESTAMP (--- contd ---)
```

```
  MegacoMessage (2nd part)
  ...
  TRACE_TIMESTAMP (--- contd ---)
  MegacoMessage (n'th part)
  [END #SEQNUM]
```

Here is an example of a divided message:

```
Oct 10 09:32:56
  [BEGIN #779  Oct 10 09:32:56  UDP:22.0.0.6:2944->10.50.40.100:2944  RX  436]
  AU=0x00000000:0x00000000:0x00000000000000000000000000
  MEGACO/3 [22.0.0.6]:2944
  Transaction = 12567418{
      Context = 65323 {
         Modify=ip/4/vif-0/3 {
            Media{
               Stream=1{
                  LocalControl{Mode=SR,DS/DSCP=00,TMAN/POL=on}}}},
         Modify=ip/4/vif-0/3 {
            Media{
               Stream=1{
Oct 10 14:13:19 (--- contd ---)
                  LocalControl{Mode=SR,DS/DSCP=00,TMAN/POL=on}}}}}}
  [END #779]
```

## *Configuring Logging of H.248 Messages*

The captured data is stored by default in the pgcpd file in the /var/log directory.

**Step-by-Step Procedure** To enable logging of H.248 messages:

1. Access the PGCP traceoptions configuration.

   ```
   [edit]
   user@host#edit services pgcp traceoptions
   ```

2. Set the flag that enables logging of H.248 messages.

   ```
   [edit services pgcp traceoptions]
   user@host#set flag pgcp-stack-h248
   ```

3. Set the maximum number of trace files. If you specify a maximum number of files, you also must specify a maximum file size.

   ```
   [edit services pgcp traceoptions]
   user@host#set file files 10
   ```

4. Set the maximum size of the trace files.

   ```
   [edit services pgcp traceoptions]
   user@host#set file pgcpd size 10000k
   ```

**Related Topics**

- *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*

Chapter 9

# Example: Providing Voice Solutions in a Next-Generation Network

This example describes how to configure the voice solution on a router in the service provider core network. Topics include:



## Requirements

This example uses the following software and hardware components:

- JUNOS Release 9.0 or later
- T640 routing node with a MultiServices 500 PIC

## Overview and Topology

This example shows how to configure the SP PG router in the topology shown in Figure 25 on page 136.

**Figure 25: Voice Solution Topology Diagram**

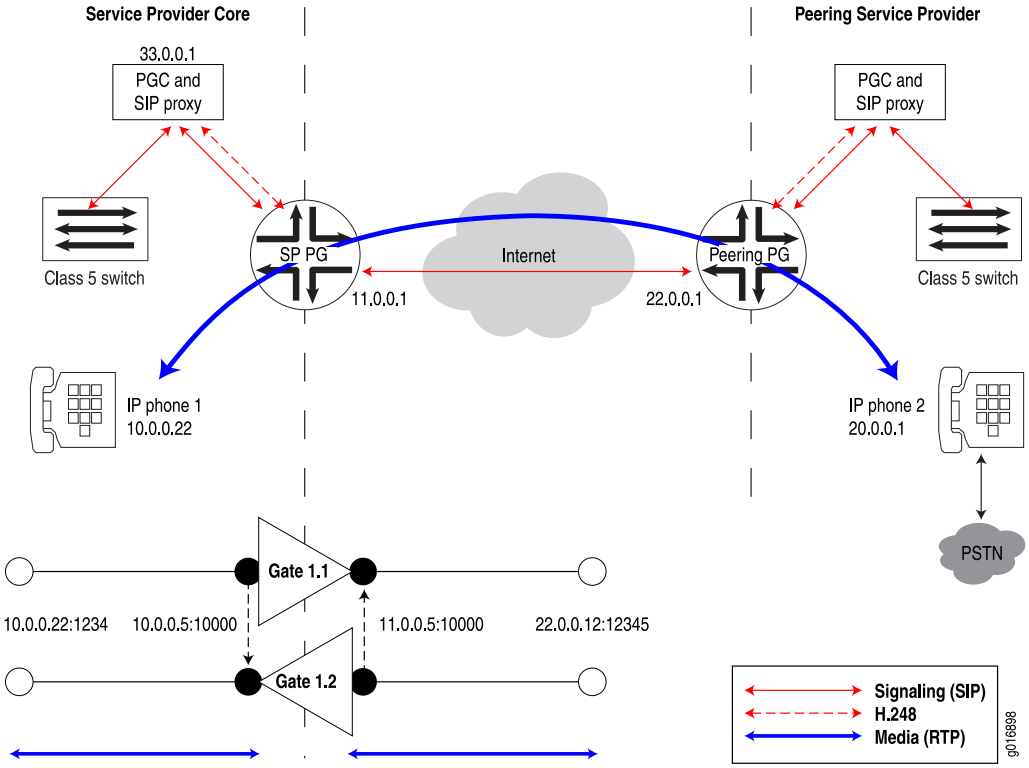


Table 17 on page 137 shows the voice configuration components.

**Table 17: Addresses Used in the Voice Solution Topology**

| Device | Interfaces/VPG/NAT Pools | Address |
|---|---|---|
| SP PG (Service Provider Packet Gateway) | fe-1/1/0 | 172.16.10.1<br>172.16.20.2<br>172.16.30.3 |
| | fe-1/1/1 | 11.0.0.1<br>2001:db8:12:3::1/112 |
| | fe-1/1/2 | 10.0.0.1<br>2001:db8:14:4::1/112 |
| | sp–1/2/0.10 (inside service domain for VPG-1)<br>sp-1/2/0.20 (outside service domain for VPG-1)<br>sp-1/2/0.30 (inside service domain for VPG-2)<br>sp-1/2/0.40 (outside service domain for VPG-2) | Not applicable |
| | sp-1/2/1.10 (inside service domain for VPG-3)<br>sp-1/2/1.20 (outside service domain for VPG 3) | Not applicable |
| | VPG-1—Provides both SIP and RTP over IPv4<br>VPG-2—Provides RTP (video) over IPv4<br>VPG-3—Provides media over IPv6 | 172.16.10.1<br>172.16.20.2<br>172.16.30.3 |
| | Media (RTP) NAT Pools<br>■ vpg1_peer_rtp-nat-pool-1<br>■ vpg2_peer_rtp-nat-pool-2<br>■ vpg3_peer_rtp-nat-pool-3<br>■ vpg1_core_rtp-nat-pool-4<br>■ vpg2_core_rtp-nat-pool-5<br>■ vpg3_core_rtp-nat-pool-6 | Pool Address<br>■ 11.0.0.5<br>■ 11.0.0.25<br>■ 2001:db8:10:3::100/128<br>■ 10.0.0.5<br>■ 10.0.0.25<br>■ 2001:db8:13:2::100/128 |
| | Signaling (SIP) NAT Pools<br>■ vpg1_peer_sip-nat-pool-7<br>■ vpg1_core_sip-nat-pool-8 | Pool Address<br>■ 11.0.0.2<br>■ 10.0.0.2 |
| Packet Gateway Controller | | 33.0.0.1 |
| Peering Router | | 22.0.0.1 |

## Configuration

To configure the SP PG router:



**CLI Quick Configuration** To quickly configure this example on the SP PG router, copy the following commands, and paste them into the router terminal window:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 172.16.10.1
set interfaces fe-1/1/0 unit 0 family inet address 172.16.20.2
set interfaces fe-1/1/0 unit 0 family inet address 172.16.30.3
set interfaces fe-1/1/1 description peer
set interfaces fe-1/1/1 unit 0 family inet address 11.0.0.1
set interfaces fe-1/1/1 unit 0 family inet6 address 2001:db8:12:3::1/112
set interfaces fe-1/1/2 description core
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.1
set interfaces fe-1/1/2 unit 0 family inet6 address 2001:db8:14:4::1/112
set interfaces sp-1/2/0 description vpg1_vpg2_pgcp_service_ipv4
set interfaces sp-1/2/0 unit 0 family inet
set interfaces sp-1/2/0 unit 10 family inet
set interfaces sp-1/2/0 unit 10 service-domain inside
set interfaces sp-1/2/0 unit 20 family inet
set interfaces sp-1/2/0 unit 20 service-domain outside
set interfaces sp-1/2/0 unit 30 family inet
set interfaces sp-1/2/0 unit 30 service-domain inside
set interfaces sp-1/2/0 unit 40 family inet
set interfaces sp-1/2/0 unit 40 service-domain outside
set interfaces sp-1/2/0 traceoptions flag all
set interfaces sp-1/2/0 services-options syslog host local services any
set interfaces sp-1/2/1 description vpg3_pgcp_service_ipv6
set interfaces sp-1/2/1 unit 0 family inet6
set interfaces sp-1/2/1 unit 10 family inet6
set interfaces sp-1/2/1 unit 10 service-domain inside
set interfaces sp-1/2/1 unit 20 family inet6
set interfaces sp-1/2/1 unit 20 service-domain outside
set interfaces sp-1/2/1 traceoptions flag all
set interfaces sp-1/2/1 services-options syslog host local services any
set services pgcp gateway vpg-1 gateway-address 172.16.10.1
set services pgcp gateway vpg-1 gateway-port 2944
set services pgcp gateway vpg-1 cleanup-timeout 3600
```

```
set services pgcp gateway vpg-2 gateway-address 172.16.20.2
set services pgcp gateway vpg-2 gateway-port 2944
set services pgcp gateway vpg-2 cleanup-timeout 3600
set services pgcp gateway vpg-3 gateway-address 172.16.30.3
set services pgcp gateway vpg-3 gateway-port 2944
set services pgcp gateway vpg-3 cleanup-timeout 3600
set services pgcp gateway vpg-1 gateway-controller pgc-1 controller-address 33.0.0.1
set services pgcp gateway vpg-1 gateway-controller pgc-1 controller-port 2944
set services pgcp gateway vpg-1 gateway-controller pgc-1 interim-ah-scheme
  algorithm hmac-null
set services pgcp gateway vpg-2 gateway-controller pgc-1 controller-address 33.0.0.1
set services pgcp gateway vpg-2 gateway-controller pgc-1 controller-port 2944
set services pgcp gateway vpg-2 gateway-controller pgc-1 interim-ah-scheme
  algorithm hmac-null
set services pgcp gateway vpg-3 gateway-controller pgc-1 controller-address 33.0.0.1
set services pgcp gateway vpg-3 gateway-controller pgc-1 controller-port 2944
set services pgcp gateway vpg-3 gateway-controller pgc-1 interim-ah-scheme
  algorithm hmac-null
set services nat pool vpg1_peer_rtp-nat-pool-1 address 11.0.0.5
set services nat pool vpg1_peer_rtp-nat-pool-1 port automatic
set services nat pool vpg1_peer_rtp-nat-pool-1 pgcp
set services nat pool vpg1_peer_rtp-nat-pool-1 pgcp ports-per-session 2
set services nat pool vpg2_peer_rtp-nat-pool-2 address 11.0.0.25
set services nat pool vpg2_peer_rtp-nat-pool-2 port automatic
set services nat pool vpg2_peer_rtp-nat-pool-2 pgcp
set services nat pool vpg2_peer_rtp-nat-pool-2 pgcp ports-per-session 2
set services nat pool vpg3_peer_rtp-nat-pool-3 address 2001:db8:10:3::100/128
set services nat pool vpg3_peer_rtp-nat-pool-3 port automatic
set services nat pool vpg3_peer_rtp-nat-pool-3 pgcp
set services nat pool vpg3_peer_rtp-nat-pool-3 pgcp ports-per-session 2
set services nat pool vpg1_core_rtp-nat-pool-4 address 10.0.0.5
set services nat pool vpg1_core_rtp-nat-pool-4 port automatic
set services nat pool vpg1_core_rtp-nat-pool-4 pgcp
set services nat pool vpg1_core_rtp-nat-pool-4 pgcp ports-per-session 2
set services nat pool vpg2_core_rtp-nat-pool-5 address 10.0.0.25
set services nat pool vpg2_core_rtp-nat-pool-5 port automatic
set services nat pool vpg2_core_rtp-nat-pool-5 pgcp
set services nat pool vpg2_core_rtp-nat-pool-5 pgcp ports-per-session 4
set services nat pool vpg3_core_rtp-nat-pool-6 address 2001:db8:13:2::100/128
set services nat pool vpg3_core_rtp-nat-pool-6 port automatic
set services nat pool vpg3_core_rtp-nat-pool-6 pgcp
set services nat pool vpg3_core_rtp-nat-pool-6 pgcp ports-per-session 2
set services nat pool vgp1_peer_sip-nat-pool-7 address 11.0.0.2
set services nat pool vgp1_peer_sip-nat-pool-7 port range low 10000 high 50000
set services nat pool vgp1_peer_sip-nat-pool-7 pgcp
set services nat pool vgp1_peer_sip-nat-pool-7 pgcp remotely-controlled
set services nat pool vgp1_peer_sip-nat-pool-7 pgcp ports-per-session 2
set services nat pool vpg1_core_sip-nat-pool-8 address 10.0.0.2
set services nat pool vpg1_core_sip-nat-pool-8 port range low 10000 high 50000
set services nat pool vpg1_core_sip-nat-pool-8 pgcp
set services nat pool vpg1_core_sip-nat-pool-8 pgcp remotely-controlled
set services nat pool vpg1_core_sip-nat-pool-8 pgcp ports-per-session 2
set services pgcp media-service vpg1_peer_rtp_ms1 nat-pool vpg1_peer_rtp-nat-pool-1
set services pgcp media-service vpg1_core_rtp_ms4 nat-pool vpg1_core_rtp-nat-pool-4
set services pgcp media-service vpg1_peer_sip_ms7 nat-pool
  vgp1_peer_sip-nat-pool-7
```

```
set services pgcp media-service vpg1_core_sip_ms8 nat-pool
  vpg1_core_sip-nat-pool-8
set services pgcp media-service vpg2_peer_rtp_ms2 nat-pool vpg2_peer_rtp-nat-pool-2
set services pgcp media-service vpg2_core_rtp_ms5 nat-pool vpg2_core_rtp-nat-pool-5
set services pgcp media-service vpg3_peer_rtp_ms3 nat-pool vpg3_peer_rtp-nat-pool-3
set services pgcp media-service vpg3_core_rtp_ms6 nat-pool vpg3_core_rtp-nat-pool-6
set services pgcp virtual-interface v1 media-service vpg1_core_rtp_ms4
set services pgcp virtual-interface v1 media-service vpg1_core_sip_ms8
set services pgcp virtual-interface v1 media-service vpg2_core_rtp_ms5
set services pgcp virtual-interface v1 media-service vpg3_core_rtp_ms6
set services pgcp virtual-interface v1 interface fe-1/1/1
set services pgcp virtual-interface v2 media-service vpg1_peer_rtp_ms1
set services pgcp virtual-interface v2 media-service vpg1_peer_sip_ms7
set services pgcp virtual-interface v2 media-service vpg2_peer_rtp_ms2
set services pgcp virtual-interface v2 media-service vpg3_peer_rtp_ms3
set services pgcp virtual-interface v2 interface fe-1/1/2
set services pgcp rule pgcp-rule1 gateway vpg-1 media-service vpg1_peer_rtp_ms1
set services pgcp rule pgcp-rule1 gateway vpg-1 media-service vpg1_peer_sip_ms7
set services pgcp rule pgcp-rule1 gateway vpg-1 media-service vpg1_core_rtp_ms4
set services pgcp rule pgcp-rule1 gateway vpg-1 media-service vpg1_core_sip_ms8
set services pgcp rule pgcp-rule2 gateway vpg-2 media-service vpg2_peer_rtp_ms2
set services pgcp rule pgcp-rule2 gateway vpg-2 media-service vpg2_core_rtp_ms5
set services pgcp rule pgcp-rule3 gateway vpg-3 media-service vpg3_peer_rtp_ms3
set services pgcp rule pgcp-rule3 gateway vpg-3 media-service vpg3_core_rtp_ms6
set services stateful-firewall rule r1 match-direction input-output term t1 then reject
set services service-set vpg1_pgcp-svc-set pgcp-rules pgcp-rule1
set services service-set vpg1_pgcp-svc-set stateful-firewall-rules r1
set services service-set vpg1_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/0.10
set services service-set vpg1_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/0.20
set services service-set vpg1_pgcp-svc-set syslog host local-1 services any
set services service-set vpg2_pgcp-svc-set pgcp-rules pgcp-rule2
set services service-set vpg2_pgcp-svc-set stateful-firewall-rules r1
set services service-set vpg2_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/0.30
set services service-set vpg2_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/0.40
set services service-set vpg2_pgcp-svc-set syslog host local-1 services any
set services service-set vpg3_pgcp-svc-set pgcp-rules pgcp-rule3
set services service-set vpg3_pgcp-svc-set stateful-firewall-rules r1
set services service-set vpg3_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/1.10
set services service-set vpg3_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/1.20
set services service-set vpg3_pgcp-svc-set syslog host local-1 services any
set services pgcp gateway vpg-1 h248-properties diffserv dscp 1D
```

### *Configuring Physical Interfaces*

**CLI Quick Configuration** To quickly configure physical interfaces on the SP PG router, copy the following commands and paste them into the router terminal window:

```
[edit interfaces]
set fe-1/1/0 unit 0 family inet address 172.16.10.1
```

```
set fe-1/1/0 unit 0 family inet address 172.16.20.2
set fe-1/1/0 unit 0 family inet address 172.16.30.3
set fe-1/1/1 description peer
set fe-1/1/1 unit 0 family inet address 11.0.0.1
set fe-1/1/1 unit 0 family inet6 address 2001:db8:12:3::1/112
set fe-1/1/2 description core
set fe-1/1/2 unit 0 family inet address 10.0.0.1
set fe-1/1/2 unit 0 family inet6 address 2001:db8:14:4::1/112
```

**Step-by-Step Procedure** To configure physical interfaces on the SP PG router:

1. Configure an interface with the addresses of the VPGs.

```
[edit interfaces]
user@sp-pg-router# edit fe-1/1/0 unit 0 family inet

[edit interfaces fe-1/1/0 unit 0 family inet]
user@sp-pg-router# set address 172.16.10.1
user@sp-pg-router# set address 172.16.20.2
user@sp-pg-router# set address 172.16.30.3
```

2. Configure an interface for the connection to the peering router.

```
[edit interfaces]
user@sp-pg-router# edit fe-1/1/1

[edit interfaces fe-1/1/1]
user@sp-pg-router# set description peer
user@sp-pg-router# set unit 0 family inet address 11.0.0.1
user@sp-pg-router# set unit 0 family inet6 address 2001:db8:12:3::1/112
```

3. Configure an interface for the connection to the IP phone.

```
[edit interfaces]
user@sp-pg-router# edit fe-1/1/2

[edit interfaces fe-1/1/2]
user@sp-pg-router# set description core
user@sp-pg-router# set unit 0 family inet address 10.0.0.1
user@sp-pg-router# set unit 0 family inet6 address 2001:db8:14:4::1/112
```

**Configuration Results** Display the results of the configuration.

```
[edit interfaces]
user@sp-pg-router# show
fe-1/1/0 {
    unit 0 {
        family inet {
            address 172.16.10.1/32;
            address 172.16.20.2/32;
            address 172.16.30.3/32;
        }
    }
}
    fe-1/1/1 {
```

```
    description peer;
    unit 0 {
      family inet {
        address 11.0.0.1/32;
      }
      family inet6 {
        address 2001:db8:12:3::1/112;
      }
    }
  }
  fe-1/1/2 {
    description core;
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
      family inet6 {
        address 2001:db8:14:4::1/112;
      }
    }
  }
```

### *Configuring the Service Interfaces*

**CLI Quick Configuration** To quickly configure the service interfaces, copy the following commands and paste them into the router terminal window:

```
[edit interfaces]
set sp-1/2/0 description vpg1_vpg2_pgcp_service_ipv4
set sp-1/2/0 unit 0 family inet
set sp-1/2/0 unit 10 family inet
set sp-1/2/0 unit 10 service-domain inside
set sp-1/2/0 unit 20 family inet
set sp-1/2/0 unit 20 service-domain outside
set sp-1/2/0 unit 30 family inet
set sp-1/2/0 unit 30 service-domain inside
set sp-1/2/0 unit 40 family inet
set sp-1/2/0 unit 40 service-domain outside
set sp-1/2/0 traceoptions flag all
set sp-1/2/0 services-options syslog host local services any
set sp-1/2/1 description vpg3_pgcp_service_ipv6
set sp-1/2/1 unit 0 family inet6
set sp-1/2/1 unit 10 family inet6
set sp-1/2/1 unit 10 service-domain inside
set sp-1/2/1 unit 20 family inet6
set sp-1/2/1 unit 20 service-domain outside
set sp-1/2/1 traceoptions flag all
set sp-1/2/1 services-options syslog host local services any
```

**Step-by-Step Procedure** To configure the service interface:

1. Create the interface, and enter edit mode for the interface.

```
[edit interfaces]
user@sp-pg-router#edit sp-1/2/0
```

2. Configure an IPv4 service interface for VPG–1 and VPG–2.

```
[edit interfaces sp-1/2/0]
user@sp-pg-router#set description vpg1_vpg2_pgcp_service_ipv4
user@sp-pg-router#set unit 0 family inet
user@sp-pg-router#set unit 10 family inet
user@sp-pg-router#set unit 10 service-domain inside
user@sp-pg-router#set unit 20 family inet
user@sp-pg-router#set unit 20 service-domain outside
user@sp-pg-router#set unit 30 family inet
user@sp-pg-router#set unit 30 service-domain inside
user@sp-pg-router#set unit 40 family inet
user@sp-pg-router#set unit 40 service-domain outside
user@sp-pg-router#set traceoptions flag all
user@sp-pg-router#set services-options syslog host local services any
```

3. Configure an IPv6 service interface for VPG–3.

```
[edit interfaces sp-1/2/1]
user@sp-pg-router#set sp-1/2/1 description vpg3_pgcp_service_ipv6
user@sp-pg-router#set sp-1/2/1 unit 0 family inet6
user@sp-pg-router#set sp-1/2/1 unit 10 family inet6
user@sp-pg-router#set sp-1/2/1 unit 10 service-domain inside
user@sp-pg-router#set sp-1/2/1 unit 20 family inet6
user@sp-pg-router#set sp-1/2/1 unit 20 service-domain outside
user@sp-pg-router#set sp-1/2/1 traceoptions flag all
user@sp-pg-router#set sp-1/2/1 services-options syslog host local services any
```

**Configuration Results** Display the results of the configuration.

```
[edit interfaces]
user@sp-pg-router# show
sp-1/2/0 {
    description vpg1_vpg2_pgcp_service_ipv4;
    traceoptions {
        flag all;
    }
    services-options {
        syslog {
            host local {
                services any;
            }
        }
    }
    unit 0 {
        family inet;
    }
    unit 10 {
        family inet;
        service-domain inside;
    }
    unit 20 {
        family inet;
        service-domain outside;
    }
```

```
    unit 30 {
      family inet;
      service-domain inside;
    }
    unit 40 {
      family inet;
      service-domain outside;
    }
  }
  sp-1/2/1 {
    description vpg3_pgcp_service_ipv6;
    traceoptions {
      flag all;
    }
    services-options {
      syslog {
        host local {
          services any;
        }
      }
    }
    unit 0 {
      family inet6;
    }
    unit 10 {
      family inet6;
      service-domain inside;
    }
    unit 20 {
      family inet6;
      service-domain outside;
    }
  }
```

### *Configuring the Virtual Packet Gateways*

**CLI Quick Configuration** To quickly configure the VPGs, copy the following commands and paste them into the router terminal window:

```
[edit services pgcp]
set gateway vpg-1 gateway-address 172.16.10.1
set gateway vpg-1 gateway-port 2944
set gateway vpg-1 cleanup-timeout 3600
set gateway vpg-1 gateway-controller pgc-1 controller-address 33.0.0.1
set gateway vpg-1 gateway-controller pgc-1 controller-port 2944
set gateway vpg-1 gateway-controller pgc-1 interim-ah-scheme algorithm hmac-null
set gateway vpg-2 gateway-address 172.16.20.2
set gateway vpg-2 gateway-port 2944
set gateway vpg-2 cleanup-timeout 3600
set gateway vpg-2 gateway-controller pgc-1 controller-address 33.0.0.1
set gateway vpg-2 gateway-controller pgc-1 controller-port 2944
set gateway vpg-2 gateway-controller pgc-1 interim-ah-scheme algorithm hmac-null
set gateway vpg-3 gateway-address 172.16.30.3
set gateway vpg-3 gateway-port 2944
set gateway vpg-3 cleanup-timeout 3600
```

```
set gateway vpg-3 gateway-controller pgc-1 controller-address 33.0.0.1
set gateway vpg-3 gateway-controller pgc-1 controller-port 2944
set gateway vpg-3 gateway-controller pgc-1 interim-ah-scheme algorithm hmac-null
```

**Step-by-Step Procedure** To configure the VPGs:

1. Configure VPG-1.

```
[edit services pgcp]
user@sp-pg-router#edit gateway vpg-1

[edit services pgcp gateway vpg-1]
user@sp-pg-router#set gateway-address 172.16.10.1
user@sp-pg-router#set gateway-port 2944
user@sp-pg-router#set cleanup-timeout 3600
user@sp-pg-router#set gateway-controller pgc-1 controller-address 33.0.0.1
user@sp-pg-router#set gateway-controller pgc-1 controller-port 2944
user@sp-pg-router#set gateway-controller pgc-1 interim-ah-scheme algorithm
  hmac-null
```

2. Configure VPG-2.

```
[edit services pgcp]
user@sp-pg-router#edit gateway vpg-2

[edit services pgcp gateway vpg-2]
user@sp-pg-router#set gateway-address 172.16.20.2
user@sp-pg-router#set gateway-port 2944
user@sp-pg-router#set cleanup-timeout 3600
user@sp-pg-router#set gateway-controller pgc-1 controller-address 33.0.0.1
user@sp-pg-router#set gateway-controller pgc-1 controller-port 2944
user@sp-pg-router#set gateway-controller pgc-1 interim-ah-scheme algorithm
  hmac-null
```

3. Configure VPG-3.

```
[edit services pgcp]
user@sp-pg-router#edit gateway vpg-3

[edit services pgcp gateway vpg-3]
user@sp-pg-router#set gateway-address 172.16.30.3
user@sp-pg-router#set gateway-port 2944
user@sp-pg-router#set cleanup-timeout 3600
user@sp-pg-router#set gateway-controller pgc-1 controller-address 33.0.0.1
user@sp-pg-router#set gateway-controller pgc-1 controller-port 2944
user@sp-pg-router#set gateway-controller pgc-1 interim-ah-scheme algorithm
  hmac-null
```

**Configuration Results** Display the results of the configuration.

```
[edit services pgcp]
user@sp-pg-router# show
gateway vpg-1 {
  gateway-address 172.16.10.1;
  gateway-port 2944;
```

```
    cleanup-timeout 3600;
    gateway-controller pgc-1 {
      controller-address 33.0.0.1;
      controller-port 2944;
      interim-ah-scheme {
        algorithm hmac-null;
      }
    }
    }
    gateway vpg-2 {
      gateway-address 172.16.20.2;
      gateway-port 2944;
      cleanup-timeout 3600;
      gateway-controller pgc-1 {
        controller-address 33.0.0.1;
        controller-port 2944;
        interim-ah-scheme {
          algorithm hmac-null;
        }
      }
    }
    gateway vpg-3 {
      gateway-address 172.16.30.3;
      gateway-port 2944;
      cleanup-timeout 3600;
      gateway-controller pgc-1 {
        controller-address 33.0.0.1;
        controller-port 2944;
        interim-ah-scheme {
          algorithm hmac-null;
        }
      }
    }
  ## Warning: missing mandatory statement(s): 'virtual-interface'
```

### *Configuring NAT Pools for the Packet Gateway*

**CLI Quick Configuration** To quickly configure the NAT pools, copy the following commands and paste them into the router terminal window:

```
[edit services nat]
set pool vpg1_peer_rtp-nat-pool-1 address 11.0.0.5
set pool vpg1_peer_rtp-nat-pool-1 port automatic
set pool vpg1_peer_rtp-nat-pool-1 pgcp
set pool vpg1_peer_rtp-nat-pool-1 pgcp ports-per-session 2
set pool vpg2_peer_rtp-nat-pool-2 address 11.0.0.25
set pool vpg2_peer_rtp-nat-pool-2 port automatic
set pool vpg2_peer_rtp-nat-pool-2 pgcp
set pool vpg2_peer_rtp-nat-pool-2 pgcp ports-per-session 2
set pool vpg3_peer_rtp-nat-pool-3 address 2001:db8:10:3::100/128
set pool vpg3_peer_rtp-nat-pool-3 port automatic
set pool vpg3_peer_rtp-nat-pool-3 pgcp
set pool vpg3_peer_rtp-nat-pool-3 pgcp ports-per-session 2
set pool vpg1_core_rtp-nat-pool-4 address 10.0.0.5
set pool vpg1_core_rtp-nat-pool-4 port automatic
```

```
set pool vpg1_core_rtp-nat-pool-4 pgcp
set pool vpg1_core_rtp-nat-pool-4 pgcp ports-per-session 2
set pool vpg2_core_rtp-nat-pool-5 address 10.0.0.25
set pool vpg2_core_rtp-nat-pool-5 port automatic
set pool vpg2_core_rtp-nat-pool-5 pgcp
set pool vpg2_core_rtp-nat-pool-5 pgcp ports-per-session 4
set pool vpg3_core_rtp-nat-pool-6 address 2001:db8:13:2::100/128
set pool vpg3_core_rtp-nat-pool-6 port automatic
set pool vpg3_core_rtp-nat-pool-6 pgcp
set pool vpg3_core_rtp-nat-pool-6 pgcp ports-per-session 2
set pool vgp1_peer_sip-nat-pool-7 address 11.0.0.2
set pool vgp1_peer_sip-nat-pool-7 port range low 10000 high 50000
set pool vgp1_peer_sip-nat-pool-7 pgcp
set pool vgp1_peer_sip-nat-pool-7 pgcp remotely-controlled
set pool vgp1_peer_sip-nat-pool-7 pgcp ports-per-session 2
set pool vpg1_core_sip-nat-pool-8 address 10.0.0.2
set pool vpg1_core_sip-nat-pool-8 port range low 10000 high 50000
set pool vpg1_core_sip-nat-pool-8 pgcp
set pool vpg1_core_sip-nat-pool-8 pgcp remotely-controlled
set pool vpg1_core_sip-nat-pool-8 pgcp ports-per-session 2
```

**Step-by-Step Procedure** To configure NAT pools:

1. Create a media (RTP) NAT pool for VPG-1 for the access (peering) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg1_peer_rtp-nat-pool-1

[edit services nat pool vpg1_peer_rtp-nat-pool-1]
user@sp-pg-router#set address 11.0.0.5
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set pool pg1_peer_rtp-nat-pool-1 pgcp ports-per-session 2
```

2. Create a media (RTP) NAT pool for VPG-2 for the access (peering) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg2_peer_rtp-nat-pool-2

[edit services nat pool vpg2_peer_rtp-nat-pool-2]
user@sp-pg-router#set address 11.0.0.25
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp ports-per-session 2
```

3. Create a media (RTP) NAT pool for VPG-3 for the access (peering) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg3_peer_rtp-nat-pool-3

[edit services nat pool vpg3_peer_rtp-nat-pool-3]
user@sp-pg-router#set address 2001:db8:10:3::100/128
```

```
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp ports-per-session 2
```

4. Create a media (RTP) NAT pool for VPG-1 for the backbone (service provider) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg1_core_rtp-nat-pool-4

[edit services nat pool vpg1_core_rtp-nat-pool-4]
user@sp-pg-router#set address 10.0.0.5
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp ports-per-session 2
```

5. Create a media (RTP) NAT pool for VPG-2 for the backbone (service provider) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg2_core_rtp-nat-pool-5

[edit services nat pool vpg2_core_rtp-nat-pool-5]
user@sp-pg-router#set address 10.0.0.25
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set ports-per-session 4
```

6. Create a media (RTP) NAT pool for VPG-3 for the backbone (service provider) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg3_core_rtp-nat-pool-6

[edit services nat pool vpg3_core_rtp-nat-pool-6]
user@sp-pg-router#set address 2001:db8:13:2::100/128
user@sp-pg-router#set port automatic
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp ports-per-session 2
```

7. Configure a signaling (SIP) NAT pool for VPG-1 for the access (peering) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vgp1_peer_sip-nat-pool-7

[edit services nat pool vgp1_peer_sip-nat-pool-7]
user@sp-pg-router#set address 11.0.0.2
user@sp-pg-router#set port range low 10000 high 50000
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp remotely-controlled
user@sp-pg-router#set ports-per-session 2
```

8. Configure a signaling (SIP) NAT pool for the backbone (service provider) side of the network.

```
[edit services nat]
user@sp-pg-router#edit pool vpg1_core_sip-nat-pool-8

[edit services nat pool vpg1_core_sip-nat-pool-8]
user@sp-pg-router#set address 10.0.0.2
user@sp-pg-router#set port range low 10000 high 50000
user@sp-pg-router#set pgcp
user@sp-pg-router#set pgcp remotely-controlled
user@sp-pg-router#set pgcp ports-per-session 2
```

**Configuration Results** Display the results of the configuration.

```
[edit services nat]
user@sp-pg-router# show
pool vpg1_peer_rtp-nat-pool-1 {
    pgcp {
        ports-per-session 2;
    }
    address 11.0.0.5/32;
    port automatic;
}
    pool vpg2_peer_rtp-nat-pool-2 {
        pgcp {
            ports-per-session 2;
        }
        address 11.0.0.25/32;
        port automatic;
    }
    pool vpg3_peer_rtp-nat-pool-3 {
        pgcp {
            ports-per-session 2;
        }
        address 2001:db8:10:3::100/128;
        port automatic;
    }
    pool vpg1_core_rtp-nat-pool-4 {
        pgcp {
            ports-per-session 2;
        }
        address 10.0.0.5/32;
        port automatic;
    }
    pool vpg2_core_rtp-nat-pool-5 {
        pgcp {
            ports-per-session 4;
        }
        address 10.0.0.25/32;
        port automatic;
    }
    pool vpg3_core_rtp-nat-pool-6 {
        pgcp {
            ports-per-session 2;
        }
        address 2001:db8:13:2::100/128;
        port automatic;
```

```
}
pool vgp1_peer_sip-nat-pool-7 {
   pgcp {
      remotely-controlled;
      ports-per-session 2;
   }
   address 11.0.0.2/32;
   port range low 10000 high 50000;
}
pool vpg1_core_sip-nat-pool-8 {
   pgcp {
      remotely-controlled;
      ports-per-session 2;
   }
   address 10.0.0.2/32;
   port range low 10000 high 50000;
}
```

### *Assigning the NAT Pools to a Media Service*

**CLI Quick Configuration** To quickly create media services and assign NAT pools to a media service, copy the following commands and paste them into the router terminal window:

```
[edit services pgcp]
set media-service vpg1_peer_rtp_ms1 nat-pool vpg1_peer_rtp-nat-pool-1
set media-service vpg1_core_rtp_ms4 nat-pool vpg1_core_rtp-nat-pool-4
set media-service vpg1_peer_sip_ms7 nat-pool vgp1_peer_sip-nat-pool-7
set media-service vpg1_core_sip_ms8 nat-pool vpg1_core_sip-nat-pool-8
set media-service vpg2_peer_rtp_ms2 nat-pool vpg2_peer_rtp-nat-pool-2
set media-service vpg2_core_rtp_ms5 nat-pool vpg2_core_rtp-nat-pool-5
set media-service vpg3_peer_rtp_ms3 nat-pool vpg3_peer_rtp-nat-pool-3
set media-service vpg3_core_rtp_ms6 nat-pool vpg3_core_rtp-nat-pool-6
```

**Step-by-Step Procedure** To configure a media service:

1. Configure media services for each of the NAT pools for VPG-1.

```
[edit services pgcp]
user@sp-pg-router#set media-service vpg1_peer_rtp_ms1 nat-pool
   vpg1_peer_rtp-nat-pool-1
user@sp-pg-router#set media-service vpg1_core_rtp_ms4 nat-pool
   vpg1_core_rtp-nat-pool-4
user@sp-pg-router#set media-service vpg1_peer_sip_ms7 nat-pool
   vgp1_peer_sip-nat-pool-7
user@sp-pg-router#set media-service vpg1_core_sip_ms8 nat-pool
   vpg1_core_sip-nat-pool-8
```

2. Configure a media service for each of the NAT pools for VPG-2.

```
[edit services pgcp]
user@sp-pg-router#set media-service vpg2_peer_rtp_ms2 nat-pool
   vpg2_peer_rtp-nat-pool-2
user@sp-pg-router#set media-service vpg2_core_rtp_ms5 nat-pool
   vpg2_core_rtp-nat-pool-5
```

3. Configure media services for each of the NAT pools for VPG-3.

```
[edit services pgcp]
user@sp-pg-router#set media-service vpg3_peer_rtp_ms3 nat-pool
   vpg3_peer_rtp-nat-pool-3
user@sp-pg-router#set media-service vpg3_core_rtp_ms6 nat-pool
   vpg3_core_rtp-nat-pool-6
```

**Configuration Results** Display the results of the configuration.

```
[edit services pgcp]
user@sp-pg-router#show
...
media-service vpg1_peer_rtp_ms1 {
   nat-pool vpg1_peer_rtp-nat-pool-1;
   }
   media-service vpg1_core_rtp_ms4 {
      nat-pool vpg1_core_rtp-nat-pool-4;
   }
   media-service vpg2_peer_rtp_ms2 {
      nat-pool vpg2_peer_rtp-nat-pool-2;
   }
   media-service vpg2_core_rtp_ms5 {
      nat-pool vpg2_core_rtp-nat-pool-5;
   }
   media-service vpg3_peer_rtp_ms3 {
      nat-pool vpg3_peer_rtp-nat-pool-3;
   }
   media-service vpg3_core_rtp_ms6 {
      nat-pool vpg3_core_rtp-nat-pool-6;
   }
   media-service vpg1_peer_sip_ms7 {
      nat-pool vgp1_peer_sip-nat-pool-7;
   }
   media-service vpg1_core_sip_ms8 {
      nat-pool vpg1_core_sip-nat-pool-8;
   }
## Warning: missing mandatory statement(s): 'virtual-interface'
```

### *Configuring the Virtual Interfaces*

**CLI Quick Configuration** To quickly configure the virtual interfaces, copy the following commands and paste them into the router terminal window:

```
[edit services pgcp]
set virtual-interface v1 media-service vpg1_core_rtp_ms4
set virtual-interface v1 media-service vpg1_core_sip_ms8
set virtual-interface v1 media-service vpg2_core_rtp_ms5
set virtual-interface v1 media-service vpg3_core_rtp_ms6
set virtual-interface v1 interface fe-1/1/1
set virtual-interface v2 media-service vpg1_peer_rtp_ms1
set virtual-interface v2 media-service vpg1_peer_sip_ms7
set virtual-interface v2 media-service vpg2_peer_rtp_ms2
set virtual-interface v2 media-service vpg3_peer_rtp_ms3
set virtual-interface v2 interface fe-1/1/2
```

**Step-by-Step Procedure** To configure a virtual interface:

1. Create a virtual interface for the backbone (service provider) side of the network. Specify the names of the media services that contains the NAT pool to be used for gates on the virtual interface that you are configuring, and specify the physical router interface.

```
[edit services pgcp]
edit virtual-interface v1

[edit services pgcp virtual-interface v1]
user@sp-pg-router#set media-service vpg1_core_rtp_ms4
user@sp-pg-router#set media-service vpg1_core_sip_ms8
user@sp-pg-router#set media-service vpg2_core_rtp_ms5
user@sp-pg-router#set media-service vpg3_core_rtp_ms6
user@sp-pg-router#set interface fe-1/1/1
```

2. Create a virtual interface for the access (peering) side of the network. Specify the names of the media services that contains the NAT pool to be used for gates on the virtual interface that you are configuring, and specify the physical router interface.

```
[edit services pgcp]
edit virtual-interface v2

[edit services pgcp virtual-interface v2]
user@sp-pg-router#set media-service vpg1_peer_rtp_ms1
user@sp-pg-router#set media-service vpg1_peer_sip_ms7
user@sp-pg-router#set media-service vpg2_peer_rtp_ms2
user@sp-pg-router#set media-service vpg3_peer_rtp_ms3
user@sp-pg-router#set interface fe-1/1/2
```

**Configuration Results** Display the results of the configuration.

```
[edit services pgcp virtual-interface v1]
user@sp-pg-router# show
virtual-interface v1 {
   media-service [ vpg1_core_rtp_ms4 vpg1_core_sip_ms8 vpg2_core_rtp_ms5
      vpg3_core_rtp_ms6 ];
   interface fe-1/1/1.0;
}
   virtual-interface v2 {
      media-service [ vpg1_peer_rtp_ms1 vpg1_peer_sip_ms7 vpg2_peer_rtp_ms2
         vpg3_peer_rtp_ms3 ];
      interface fe-1/1/2.0;
   }
```

## *Configuring Packet Gateway Rules*

You define rules that specify the NAT pool (media service) used on a specific VPG.

**CLI Quick Configuration** To quickly define the rules, copy the following commands and paste them into the router terminal window:

```
[edit services pgcp]
```

```
set rule pgcp-rule1 gateway vpg-1 media-service vpg1_peer_rtp_ms1
set rule pgcp-rule1 gateway vpg-1 media-service vpg1_peer_sip_ms7
set rule pgcp-rule1 gateway vpg-1 media-service vpg1_core_rtp_ms4
set rule pgcp-rule1 gateway vpg-1 media-service vpg1_core_sip_ms8
set rule pgcp-rule2 gateway vpg-2 media-service vpg2_peer_rtp_ms2
set rule pgcp-rule2 gateway vpg-2 media-service vpg2_core_rtp_ms5
set rule pgcp-rule3 gateway vpg-3 media-service vpg3_peer_rtp_ms3
set rule pgcp-rule3 gateway vpg-3 media-service vpg3_core_rtp_ms6
```

**Step-by-Step Procedure** To configure the packet gateway rules:

1. Create a rule for VPG 1, and specify the media services that contains the NAT pools to be used for this VPG.e.

```
[edit services pgcp]
user@sp-pg-router#edit rule pgcp-rule1

[edit services pgcp rule pgcp-rule1]
user@sp-pg-router#set gateway vpg-1
user@sp-pg-router#set media-service vpg1_peer_rtp_ms1
user@sp-pg-router#set media-service vpg1_peer_sip_ms7
user@sp-pg-router#set media-service vpg1_core_rtp_ms4
user@sp-pg-router#set media-service vpg1_core_sip_ms8
```

2. Create a rule for VPG 2, and specify the media services that contains the NAT pools to be used for this VPG.

```
[edit services pgcp]
user@sp-pg-router#edit rule pgcp-rule2

[edit services pgcp rule pgcp-rule2]
user@sp-pg-router#set gateway vpg-2
user@sp-pg-router#set media-service vpg2_peer_rtp_ms2
user@sp-pg-router#set media-service vpg2_core_rtp_ms5
```

3. Create a rule for VPG-3, and specify the media services that contains the NAT pools to be used for this VPG.

```
[edit services pgcp]
user@sp-pg-router#edit rule pgcp-rule3

[edit services pgcp rule pgcp-rule3]
user@sp-pg-router#set gateway vpg-3
set rule pgcp-rule3 gateway vpg-3 media-service vpg3_peer_rtp_ms3
set rule pgcp-rule3 gateway vpg-3 media-service vpg3_core_rtp_ms6
```

**Configuration Results** Display the results of the configuration.

```
[edit services pgcp]
user@sp-pg-router# show
. . .
rule pgcp-rule1 {
   gateway vpg-1;
```

```
    media-service [ vpg1_peer_rtp_ms1 vpg1_peer_sip_ms7 vpg1_core_rtp_ms4
      vpg1_core_sip_ms8 ];
}
    rule pgcp-rule2 {
      gateway vpg-2;
      media-service [ vpg2_peer_rtp_ms2 vpg2_core_rtp_ms5 ];
    }
    rule pgcp-rule3 {
      gateway vpg-3;
      media-service [ vpg3_peer_rtp_ms3 vpg3_core_rtp_ms6 ];
    }
```

## *Configuring a Stateful Firewall*

You define rules that specify the NAT pool (media service) used on a specific VPG.

**CLI Quick Configuration** To quickly define the rules, copy the following commands and paste them into the router terminal window:

```
[edit services stateful-firewall]
set rule r1 match-direction input-outputset rule r1 term t1 then reject
```

**Step-by-Step Procedure** To create a stateful firewall:

1. Create a stateful firewall rule.

```
[edit services stateful-firewall]
user@host#edit rule r1
```

2. Set the match direction for the rule.

```
[edit services stateful-firewall rule r1]
user@host#set match-direction input-output
```

3. Add a term to the rule with the action set to reject.

```
[edit services stateful-firewall rule r1]
user@host#set term t1 then reject
```

**Configuration Results** Display the results of the configuration.

```
[edit services stateful-firewall]
user@sp-pg-router# show
rule r1 {
    match-direction input-output;
    term t1 {
        then {
            reject;
        }
    }
}
```

### *Configuring a Service Set*

**CLI Quick Configuration** To quickly define a service set, copy the following commands and paste them into the router terminal window:

```
[edit services]
set service-set vpg1_pgcp-svc-set pgcp-rules pgcp-rule1
set service-set vpg1_pgcp-svc-set stateful-firewall-rules r1
set service-set vpg1_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/0.10
set service-set vpg1_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/0.20
set service-set vpg1_pgcp-svc-set syslog host local-1 services any
set service-set vpg2_pgcp-svc-set pgcp-rules pgcp-rule2
set service-set vpg2_pgcp-svc-set stateful-firewall-rules r1
set service-set vpg2_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/0.30
set service-set vpg2_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/0.40
set service-set vpg2_pgcp-svc-set syslog host local-1 services any
set service-set vpg3_pgcp-svc-set pgcp-rules pgcp-rule3
set service-set vpg3_pgcp-svc-set stateful-firewall-rules r1
set service-set vpg3_pgcp-svc-set next-hop-service inside-service-interface
  sp-1/2/1.10
set service-set vpg3_pgcp-svc-set next-hop-service outside-service-interface
  sp-1/2/1.20
set service-set vpg3_pgcp-svc-set syslog host local-1 services any
```

**Step-by-Step Procedure** To configure the service sets:

1. Configure a service set for VPG-1.

```
[edit services]
user@sp-pg-router#edit service-set vpg1_pgcp-svc-set

[edit services service-set vpg1_pgcp-svc-set]
user@sp-pg-router#set pgcp-rules pgcp-rule1
user@sp-pg-router#set stateful-firewall-rules r1
user@sp-pg-router#set next-hop-service inside-service-interface sp-1/2/0.10
user@sp-pg-router#set next-hop-service outside-service-interface sp-1/2/0.20
user@sp-pg-router#set syslog host local-1 services any
```

2. Configure a service set for VPG-2.

```
[edit services]
user@sp-pg-router#edit service-set vpg2_pgcp-svc-set

[edit services service-set vpg2_pgcp-svc-set]
user@sp-pg-router#set pgcp-rules pgcp-rule2
user@sp-pg-router#set stateful-firewall-rules r1
user@sp-pg-router#set next-hop-service inside-service-interface sp-1/2/0.30
user@sp-pg-router#set next-hop-service outside-service-interface sp-1/2/0.40
user@sp-pg-router#set syslog host local-1 services any
```

3. Configure a service set for VPG-3.

```
[edit services]
user@sp-pg-router#edit service-set vpg3_pgcp-svc-set

[edit services service-set vpg3_pgcp-svc-set]
user@sp-pg-router#set pgcp-rules pgcp-rule3
user@sp-pg-router#set stateful-firewall-rules r1
user@sp-pg-router#set next-hop-service inside-service-interface sp-1/2/1.10
user@sp-pg-router#set next-hop-service outside-service-interface sp-1/2/1.20
user@sp-pg-router#set syslog host local-1 services any
```

**Configuration Results** Display the results of the configuration.

```
[edit services]
user@sp-pg-router# show service-set vpg1_pgcp-svc-set
syslog {
   host local-1 {
      services any;
   }
}
stateful-firewall-rules r1;
pgcp-rules pgcp-rule1;
next-hop-service {
   inside-service-interface sp-1/2/0.10;
   outside-service-interface sp-1/2/0.20;
}

[edit services]
user@sp-pg-router# show service-set vpg2_pgcp-svc-set
syslog {
   host local-1 {
      services any;
      }
   }
stateful-firewall-rules r1;
pgcp-rules pgcp-rule2;
next-hop-service {
   inside-service-interface sp-1/2/0.30;
   outside-service-interface sp-1/2/0.40;
}

[edit services]
user@sp-pg-router# show service-set vpg3_pgcp-svc-set
syslog {
   host local-1 {
      services any;
      }
   }
stateful-firewall-rules r1;
pgcp-rules pgcp-rule3;
next-hop-service {
   inside-service-interface sp-1/2/1.10;
   outside-service-interface sp-1/2/1.20;
}
```

### *Configuring QoS for Voice Calls*

**CLI Quick Configuration** To quickly configure a default value for the Differentiated Services (DiffServ) code point (DSCP), copy the following command and paste it into the router terminal window:

```
[edit services pgcp]
set gateway vpg-1 h248-properties diffserv dscp 1D
```

**Step-by-Step Procedure** To configure default values for H.248 segmentation properties:

1. Access the configuration of the H.248 DiffServ properties.

```
[edit services pgcp gateway vpg-1]
user@sp-pg-router#edit h248-properties diffserv
```

2. Configure a value for the DSCP.

```
[edit services pgcp gateway vpg-1 h248-properties diffserv]
user@sp-pg-router#set dscp 1D
```

**Configuration Results** Display the results of the configuration.

```
[edit services pgcp]
user@sp-pg-router# show
dscp-value: '1D': must be 8 bits bit-string or hex value in the format 0xXX at '1D'
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Active PGCP Configuration on page 157
- Verifying That Gates Are Running on page 160
- Verifying PGCP Terminations on page 161
- Verifying PGCP Flows on page 162
- Verifying H.248 Parameters Set by the PGC on page 162

### *Verifying the Active PGCP Configuration*

**Purpose** Verify the active PGCP configuration that is running on your router.

**Action**

```
user@sp-pg-router> show services pgcp active-configuration
Packet gateway media service configuration:
 Media service name: vpg1_peer_rtp_ms1
  Nat pool             :  vpg1_peer_rtp-nat-pool-1

Packet gateway media service configuration:
 Media service name: vpg2_peer_rtp_ms2
  Nat pool             :  vpg2_peer_rtp-nat-pool-2

Packet gateway media service configuration:
 Media service name: vpg3_peer_rtp_ms3
  Nat pool             :  vpg3_peer_rtp-nat-pool-3
```

```
Packet gateway media service configuration:
 Media service name: vpg1_core_rtp_ms4
  Nat pool             :  vpg1_core_rtp-nat-pool-4

Packet gateway media service configuration:
 Media service name: vpg2_core_rtp_ms5
  Nat pool             :  core_rtp-nat-pool-5

Packet gateway media service configuration:
 Media service name:  vpg3_core_rtp_ms6
  Nat pool             :  vpg3_core_rtp-nat-pool-6

Packet gateway media service configuration:
 Media service name: vpg1_peer_sip_ms7
  Nat pool             :  vgp1_peer_sip-nat-pool-7

Packet gateway media service configuration:
 Media service name: vpg1_core_sip_ms8
  Nat pool             :  vpg1_core_sip-nat-pool-8

 Packet gateway virtual interface configuration:
  Virtual Interface name: v1
  Status               :  In-Service
  Interface name       :  fe-1/1/1.0
  Media Service Name   :  vpg1_core_rtp_ms4
  Media Service Name   :  vpg1_core_sip_ms8
  Media Service Name   :  vpg2_core_rtp_ms5
  Media Service Name   :  vpg2_core_rtp_ms5

 Packet gateway virtual interface configuration:
  Virtual Interface name: v2
  Status               :  In-Service
  Interface name       :  fe-1/1/2.0
  Media Service Name   :  vpg1_peer_rtp_ms1
  Media Service Name   :  vpg1_peer_sip_ms7
  Media Service Name   :  vpg2_peer_rtp_ms2
  Media Service Name   :  vpg3_peer_rtp_ms3

Packet gateway configuration:
  Name                                  :  vpg-1
  IP address                            :  172.16.10.1
  Port                                  :  2944
  Status                                :  In-Service (Registered)
  Active gateway controller             :  pcg-1
  Replication socket                    :  Connected (Syncing)
  Cleanup timeout [secs]                :  3600
  Gate inactivity delay [secs]          :  240
  Gate inactivity duration (Q-MI ) [secs]  :  86400
  Latching Deadlock duration [secs]        : 15

Packet gateway configuration:
  Name                                  :  vpg-2
  IP address                            :  172.16.20.2
  Port                                  :  2944
  Status                                :  In-Service (Disconnected)
  Active gateway controller             :
  Replication socket                    :  Connected (Syncing)
  Cleanup timeout [secs]                :  3600
  Gate inactivity delay [secs]          :  240
  Gate inactivity duration (Q-MI ) [secs]  :  86400
```

```
  Latching Deadlock duration [secs]        : 15

Packet gateway configuration:
  Name                                     :  vpg-3
  IP address                               :  172.16.30.3
  Port                                     :  2944
  Status                                   :  In-Service (Disconnected)
  Active gateway controller                :
  Replication socket                       :  Connected (Syncing)
  Cleanup timeout [secs]                   :  3600
  Gate inactivity delay [secs]             :  240
  Gate inactivity duration (Q-MI ) [secs]  :  86400
  Latching Deadlock duration [secs]        : 15

H248 timers configuration:
 Max waiting delay (MWD) [millisec]          :  10000
 Max retransmission delay (T-MAX) [millisec] :  25000
 Initial average ack delay (I-AAD) [millisec]:  1000
 Max net propagation delay (M-NPD) [millisec]:  5000

H248 options configuration:
 Wildcard response service-change        :  NO

H248 diffserv configuration:
 dscp          :  0x00

H248 segmentation                           default
 MG segmentation timer [millisec]         :  4000
 MG maximum PDU size                      :  1472 byte
 MGC segmentation timer [millisec]        :  4000
 MGC maximum PDU size                     :  1472 byte

H248 base root                              default
 Normal MG execution time [millisec]      :  500
 MG Provisional response timer [millisec] :  2000
 MG Originated pending limit              :  4
 Normal MGC execution time [millisec]     :  500
 MGC Provisional response timer [millisec]:  4000
 MGC Originated pending limit             :  4

Fast update filters:
 Maximum terms              : 20000
 Maximum term percentage    : 10

Packet gateway controller configuration:
 Controller name                    :  pgc-1
 Controller IP address              :  33.0.0.1
 Controller port                    :  2944

Packet gateway rule configuration:
 Rule name            :  pgcp-rule1
 Gateway name         :  vpg-1

Packet gateway rule configuration:
 Rule name            :  pgcp-rule2
 Gateway name         :  vpg-2

Packet gateway rule configuration:
 Rule name            :  pgcp-rule3
 Gateway name         :  vpg-3
```

```
Packet gateway service set configuration:
 Service set name     :  vpg1_pgcp-svc-set
 Service set id       :  1
 Rule name            :  pgcp-rule1

Packet gateway service set configuration:
 Service set name     :  vpg2_pgcp-svc-set
 Service set id       :  2
 Rule name            :  pgcp-rule2

Packet gateway service set configuration:
 Service set name     :  vpg3_pgcp-svc-set
 Service set id       :  3
 Rule name            :  pgcp-rule3

Packet gateway service pics status:
 Name    : sp-1/2/0
 Status  : Connected

Packet gateway service pics status:
 Name    : sp-1/2/1
 Status  : Connected

Firewall:
 Status             :  Connected
 Number of terms    :  2
 Number of filters  :  2
```

**Meaning** Use the show services pgcp active-configuration command to see your configuration and to display the current status of your virtual interfaces, VPG (Packet gateway configuration), and services PICs.

In addition, make sure that:

- At least one VPG (Packet gateway configuration) is In-Service (Registered). VPG Out-Of-Service might mean that:
  - The PGC is down.
  - The network connection between the PGC and the VPG is down, or there is a related network problem.
  - An unknown software problem exists. In this case, review the packet gateway logs for more information or contact JTAC.
- The Replication socket is Connected. If it is not, graceful Routing Engine switchover (GRES ) is not enabled.
- The Packet gateway service PIC status is Connected. If it is not, there might be a hardware problem with the PIC.
- The Firewall Status is Connected. If it is not, no connection exists to the fast update filters for rate limiting.

### *Verifying That Gates Are Running*

**Purpose** Verify that gates are running on VPG-1.

**Action**

```
user@sp-pg-router> show services pgcp gates gateway vpg-1
Packet gateway configuration:
  Name                                    :  vpg-1
  IP address                              :  172.16.10.1
  Port                                    :  2944
  Status                                  :  In-Service (Registered)

Gate information:
Gate id: 4295033088
Gate state:  active
Action: forward
Service set id: 1
Media card: sp-1/2/0
Media handler: vpg1_pgcp-svc-set
Termination-id-string: ip/4/vif-0/2

Gate id: 4295033089
Gate state:  active
Action: forward
Service set id: 1
Media card: sp-1/2/0
Media handler: vpg1_pgcp-svc-set
Termination-id-string: ip/4/vif-0/3
```

**Meaning** The show services pgcp gates command lists the gates on the VPG. It shows whether gates are active, disabled, or closed. It also shows the current action being performed on the gate—forward, add, or drop.

### *Verifying PGCP Terminations*

**Purpose** Verify the PGCP terminations on VPG-1.

**Action**

```
user@sp-pg-router> show services pgcp terminations vpg-1
Packet gateway configuration:
  Name                                    :  vpg-1
  IP address                              :  172.16.10.1
  Port                                    :  2944
  Status                                  :  In-Service (Registered)


Termination name                   State           Duration(msecs)
ip/4/vif-0/2       In-service      390288
Gate-id         Direction       State        Action
4295033088      A->B            active       forward
4295033089      B->A            active       forward


Termination name                   State           Duration(msecs)
ip/4/vif-0/3       In-service      390294
Gate-id         Direction       State        Action
4295033088      A->B            active       forward
4295033089      B->A            active       forward
```

**Meaning** The show services pgcp terminations command lists the terminations on the VPG. It shows whether the gates within a termination are active, disabled, or closed. You can use the termination names (termination IDs) to troubleshoot problems with voice calls.

### *Verifying PGCP Flows*

**Purpose** Verify the PGCP flows.

**Action**

```
user@sp-pg-router> show services pgcp flows
Interface: sp-1/2/0, Service set: pgcp-svc-set-1
Flow                                                State    Dir       Frm count
Gate id: 4295033089
UDP            20.0.0.1:0     ->    11.0.0.5:1024  Forward  I              0
    NAT source         20.0.0.1:0       ->     10.10.0.5:1024
    NAT dest       11.0.0.5:1024    ->         10.0.0.1:20002
Gate id: 4295033089
UDP            20.0.0.1:0     ->    11.0.0.5:1025  Forward  I              0
    NAT source         20.0.0.1:0       ->     10.10.0.5:1025
    NAT dest       11.0.0.5:1025    ->         10.0.0.1:20003
Gate id: 4295033088
UDP            0.0.0.0:0      ->    10.10.0.5:1024  Forward  I              0
    NAT source         0.0.0.0:0        ->     11.0.0.5:1024
    NAT dest       10.10.0.5:1024    ->          20.0.0.1:10002
Gate id: 4295033088
UDP            0.0.0.0:0      ->    10.10.0.5:1025  Forward  I              0
    NAT source         0.0.0.0:0        ->     11.0.0.5:1025
    NAT dest       10.10.0.5:1025    ->          20.0.0.1:10003
```

### *Verifying H.248 Parameters Set by the PGC*

**Purpose** Verify the H.248 parameters that are set by the PGC.

**Action**

```
user@sp-pg-router> show services pgcp terminations vpg-1 h248
Termination information:
ip/4/vif-0/2 {
      MEDIA {
        TERMINATIONSTATE { SERVICESTATES = INSERVICE },
        STREAM = 1 {
          LOCALCONTROL { MODE = SENDRECEIVE,
                         DS/DSCP = 00,
                         TMAN/MBS = 0,
                         TMAN/PDR = 0,
                         TMAN/POL = OFF,
                         TMAN/SDR = 0,
                         MGCINFO/DB = 00,
                         GM/RSB = ON,
                         GM/SAF = OFF,
                         GM/SPF = OFF,
                         GM/SPR = 0,
                         GM/ESAS = OFF,
                         GM/ESPS = OFF,
                         GM/LSP = 0 },
          LOCAL {
v=0
c=IN IP4 10.10.0.5
m=- 1024 RTP/AVP -
b=AS:0
          },
          REMOTE {
v=0
c=IN IP4 10.0.0.1
m=- 20002 RTP/AVP -
```

```
b=AS:0
          }
        }
      },
      EVENTS = 1001 { NT/QUALERT { TH = 99, STREAM = 1 } },
      SIGNALS
    }

ip/4/vif-0/3 {
      MEDIA {
        TERMINATIONSTATE { SERVICESTATES = INSERVICE },
        STREAM = 1 {
          LOCALCONTROL { MODE = SENDRECEIVE,
                         DS/DSCP = 00,
                         TMAN/MBS = 1500,
                         TMAN/PDR = 0,
                         TMAN/POL = OFF,
                         TMAN/SDR = 125000,
                         MGCINFO/DB = 00,
                         GM/RSB = ON,
                         GM/SAF = ON,
                         GM/SAM = "[20.0.0.1]",
                         GM/SPF = OFF,
                         GM/SPR = 0,
                         GM/ESAS = OFF,
                         GM/ESPS = OFF,
                         GM/LSP = 0 },
          LOCAL {
v=0
c=IN IP4 11.0.0.5
m=- 1024 RTP/AVP -
b=AS:0
          },
          REMOTE {
v=0
c=IN IP4 4.0.0.1
m=- 10002 RTP/AVP -
b=AS:0
          }
        }
      },
      SIGNALS
    }
{master}
```

**Meaning** The show services pgcp terminations *vpg-name* h248 command presents the H.248 parameters as they have been set through PGCP requests and commands from the PGC. Comparing this output with the expected PGCP requests reveals whether there was a problem with the requests and commands that the PGC sent to the VPG.

## Troubleshooting

To troubleshoot the voice configuration:

- No Audio is Reported on a Stream on page 164

### *No Audio is Reported on a Stream*

**Problem** A call completes correctly (signaling is completed), but the media (audio) stream expected to flow through the packet gateway fails.

**Solution** Locate the failed terminations and gates.

1. Acquire the relevant gate IDs termination IDs using the show services pgcp gates command.
2. Display the H.248 parameters for the termination.

```
user@sp-pg-router> show services pgcp terminations termination-prefix h248
ip/4/vif-0/2 vpg-1

Termination information:
ip/4/vif-0/2 {
      MEDIA {
        TERMINATIONSTATE { SERVICESTATES = INSERVICE },
        STREAM = 1 {
          LOCALCONTROL { MODE = SENDRECEIVE,
                         DS/DSCP = 00,
                         TMAN/MBS = 0,
                         TMAN/PDR = 0,
                         TMAN/POL = OFF,
                         TMAN/SDR = 0,
                         MGCINFO/DB = 00,
                         GM/RSB = ON,
                         GM/SAF = OFF,
                         GM/SPF = OFF,
                         GM/SPR = 0,
                         GM/ESAS = OFF,
                         GM/ESPS = OFF,
                         GM/LSP = 0 },
          LOCAL {
v=0
c=IN IP4 3.99.99.100
m=- 1024 RTP/AVP -
b=AS:0
          },
          REMOTE {
v=0
c=IN IP4 3.0.0.1
m=- 20002 RTP/AVP -
b=AS:0
          }
        }
      },
      EVENTS = 1001 { NT/QUALERT { TH = 99, STREAM = 1 } },
      SIGNALS
    }
```

3. Display information about the gate.

```
user@sp-pg-router> show services pgcp gates 4295033088

Gate information:
Gate id: 4295033088
```

```
Gate state:  active
Action: forward
Service set id: 1
Media card: sp-1/2/0
Media handler: vpg1_pgcp-svc-set

Termination-id-string: ip/4/vif-0/2
```

Using the preceding information, review and verify the following:

1. Terminations are in the In-Service state.

   If the termination is in the Out-of-Service state, no streams are allowed access to the termination. The Out-of-Service state indicates a problem with either the resources on the packet gateway or incorrect parameters requested by the PGC.

2. The termination H.248 parameters are as expected.

   Pay special attention to the LOCAL and REMOTE parameters, and make sure they are aligned with the Session Description Protocol (SDP) offered by both elements participating in the session. Also, missing or unknown values suggest a problem with the call setup initiated by the PGC.

3. Gate actions are in the Forward state.

   If one or all gates are in the Drop state, no stream is allowed to flow through it, so one-way or no audio results. If a gate is not in the Forward state, the PGC might have failed to provide a required descriptor.

   Gate actions in the Forward state, but no media is flowing (frame count is zero or not advancing), can be caused by one of the following problems:

   - Networking or routing issues, including:
     - Stream fails to reach the router. A problem exists with the network path between the stream originator and the packet gateway gates.
     - Routing Engine or PIC failure. The stream reaches the router, but the PIC fails to receive the stream. Use the show services pgcp active-configuration command to review the PIC status, and make sure that it is In-Service.
   - Hardware or element failure. The originator fails to send the stream. Use the debug tools available on the VoIP element to verify that the streams have left the element.
   - The originator is using a different source IP address than the one reported in the H.248 termination. Verify that the H.248 termination information matches the stream source and destination received by the packet gateway and the PIC. You can use the capture feature on the router to verify that the streams are received on the Packet Forwarding Engine.

**Related Topics**

- Overview of the Voice Solution on page 39
- Configuring the Voice Solution on page 61

- For a description of PGCP statements, *Chapter 27, Summary of Packet Gateway Configuration Statements* in *JUNOS Services Interfaces Configuration Guide*
- For a description of the fields in show commands, see *Chapter 25, Packet Gateway Control Protocol Operational Mode Commands* in *JUNOS System Basics and Services Command Reference*

# Part 3
# Index

■ Index on page 169

# Index

## I



## J



## L



## M



## N



## O



## P

## Q



## R



## S



## T



## V

## W