JUNIPER NETWORKS

Network Configuration Example

Configuring, Verifying, and Troubleshooting Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding

Release

10.4

Published: 2010-10-08
Revision 1

Copyright © 2010, Juniper Networks, Inc.

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Network Configuration Example Configuring Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding
Release 10.4
Copyright © 2010, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Richard Kim
Editing: Roy Spencer, Katie Smith
Illustration: Faith Bradford
Cover Design: Edmonds Design

Revision History
October 2010—R1 Junos 10.4

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 OS 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

Copyright © 2010, Juniper Networks, Inc.

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.

Copyright © 2010, Juniper Networks, Inc.

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)).

Copyright © 2010, Juniper Networks, Inc.

Copyright © 2010, Juniper Networks, Inc.

Table of Contents

Copyright © 2010, Juniper Networks, Inc.

Copyright © 2010, Juniper Networks, Inc.

Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview

This topic includes the following sections:

Fundamentals of Hub-and-Spoke VPNs in Junos OS

Junos operating system (Junos OS) is Juniper Networks' single operating system and provides the following features:

- Powerful operating system with rich IP services tool kit
- Unmatched IP dependability and security to ensure an efficient and predictable IP infrastructure
- Enhanced security and virtual private network (VPN) capabilities from Juniper Networks Firewall/IP Security (IPsec) VPN platforms, including the Secure Services Gateway (SSG) product family

This document provides an example for configuration of a multipoint interface, which is commonly used for hub-and-spoke environments. This example uses route-based VPNs from a central hub device to multiple spoke devices. Junos OS does not support a multipoint topology with policy-based VPNs.

This document is intended for network design and security engineers, as well as anyone who requires secure connectivity over public networks.

Next-Hop Tunnel Binding Overview

You can implement a hub-and-spoke VPN topology by using the route-based VPN concept in many ways. One way of implementing a hub-and-spoke VPN topology is to configure a separate secure tunnel (st0) logical unit for every spoke site. However, if a device has many associated peer devices, then the increased number of required interfaces can be a concern from the standpoint of scaling and management.

For example, the following limitation applies to:

- SSG platform – A maximum number of tunnel interfaces can be configured for the platform.
- Junos device – A maximum number of logical interface units that can be configured for the platform.

Junos OS supports the multipoint secure tunnel interfaces with the next-hop tunnel binding (NHTB) feature for easier management and scalability. NHTB enables a device to bind multiple IPsec security associations (SAs) to a single secure tunnel interface.

Copyright © 2010, Juniper Networks, Inc.

In this example:

- The secure tunnel interface operates as a point-to-point-type link by default.
- The Junos hub device is configured with st0 interface as a multipoint interface type in the st0 unit hierarchy. You need to configure multipoint interface only on the hub site; the spokes continue to use the default point-to-point mode.
- The Junos device binds multiple IPsec VPN tunnels to a single st0 interface unit.
- The Junos device links a specific destination to a specific IPsec tunnel bound to the same st0 interface, by using the following two tables:
  - Inet.0 route table
  - NHTB table
- The Junos device maps the next-hop IP address specified in the route table entry to a particular VPN tunnel specified in the NHTB table. With this feature, a single st0 interface can support many VPN tunnels.

The maximum number of IPsec tunnels is not limited by the number of st0 interfaces that you can create, but is limited by either of the following factor (whichever is lower) :

- Route table capacity
- Maximum number of dedicated IPsec tunnels allowed

To see the maximum route and tunnel capacities for your platform, refer to the relevant product data sheet.

Understanding Route-to-Tunnel Mapping

To manage the traffic among multiple IPsec VPN tunnels bound to the same st0 interface, the security device maps the next-hop gateway IP address specified in the route table to a particular IPsec tunnel name. This scenario is explained in this section with reference to Figure 1 on page 3.

The following settings are assumed for the local security device and remote VPN peer devices:

- Local security device
  - A local security device with multiple IPsec VPNs bound to a single secure tunnel (st0) interface is in subnet 10.1.1.0/24.
  - A trusted LAN interface is in subnet 10.10.10.0/24.
  - The st0.0 interface is configured in multipoint mode.
- Remote VPN peers settings
  - The remote VPN peers with fixed st0.0 interface IP addresses are all within the same 10.1.1.0/24 subnet as the local device st0 interface.
  - The remote trusted subnets are 10.20.10.0/24, 10.30.10.0/24, and 10.40.10.0/24.

Copyright © 2010, Juniper Networks, Inc.

- The st0.0 interface is configured in point-to-point mode.

Figure 1 on page 3 shows the local device routing traffic sent from 10.10.10.10 to 10.30.10.10 through the st0.0 interface and then routed the traffic through VPN2.

Figure 1: NHTB Routes and Table Entries

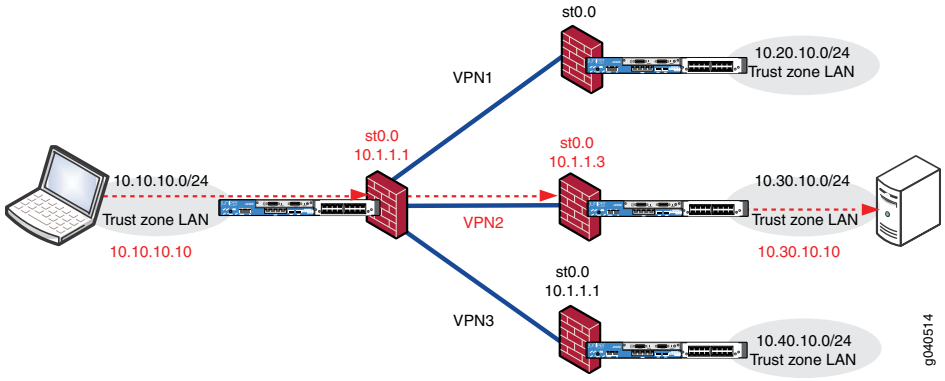

Table 1 on page 3 shows the mapping of entries in the route table to the entries in the NHTB table.

Table 1: Mapping of Route Table Entries and NHTB Table Entries

Route Table				Next-Hop Tunnel Binding Table			
IP Prefix	Next Hop	Interface		Next Hop	Interface	IPsec VPN	Flag
10.20.10.0/24	10.1.1.2	st0.0	–>	10.1.1.2	st0.0	VPN1	static
10.30.10.0/24	10.1.1.3	st0.0	–>	10.1.1.3	st0.0	VPN2	static
10.40.10.0/24	10.1.1.4	st0.0	–>	10.1.1.4	st0.0	VPN3	static

The following is an explanation of the scenario in Figure 1 on page 3 and in Table 1 on page 3:

1. You can use a dynamic routing protocol (for example, OSPF) to automatically populate the route table or you can add static routes manually. The next-hop IP address is the IP address of the remote peer's st0 interface.
2. You can use one of the following options for an NHTB table:
   - Enter the next-hop addresses manually and bind them to the appropriate IPsec VPN tunnel.

     To link the route table and NHTB table, you must enter the same IP address as the next hop, along with the appropriate IPsec VPN name in the NHTB table.
   - Allow the Junos device to acquire the next-hop address automatically from the remote peer during Phase 2 negotiations using the NOTIFY_NS_NHTB_INFORM

Copyright © 2010, Juniper Networks, Inc.

message. Note that this functionality is applicable only if both peers are Junos devices.

During Phase 2 negotiations, the two Internet Key Exchange (IKE) peers can exchange st0 interface addresses automatically through the Notify message NOTIFY_NS_NHTB_INFORM (value 40001).

3. As listed in Table 1 on page 3, the IP addresses for the next hop in the route table (which is also the same next-hop IP in the NHTB table) is the st0 interface IP address of the remote peer site. This next hop links the route –and consequently the st0 interface specified in the route -to a particular IPsec VPN tunnel.

Related Documentation

- Example: Configuring Hub-and-Spoke VPNs using Next-Hop Tunnel Binding on page 5
- Verifying Hub-and-Spoke VPN Configuration on page 37

Copyright © 2010, Juniper Networks, Inc.

Example: Configuring Hub-and-Spoke VPNs using Next-Hop Tunnel Binding

This example shows how to configure, verify, and troubleshoot the hub-and-spoke VPNs using next-hop tunnel binding.

NOTE:

- Configuration and troubleshooting details of route-based virtual private networks (VPNs) and other Junos OS specific application notes are available on the Juniper Networks Knowledge Base at http://kb.juniper.net.
- For more information on the concepts of NHTB, route-based VPNs, and interface types, refer to the complete documentation for Junos OS available at http://www.juniper.net/techpubs/.
- For more information on VPN configuration and troubleshooting, see KB10182 (http://kb.juniper.net/KB10182) available at Juniper Networks Knowledge Base.

This topic includes the following sections:

Requirements

- Junos OS Release 9.4 or later
- Junos OS with Enhanced Services Release 8.5 through 9.3
- SRX Series devices or J Series devices

Overview and Topology

Figure 2 on page 6 shows the network topology used for this example.

Copyright © 2010, Juniper Networks, Inc.

Figure 2: Network Topology

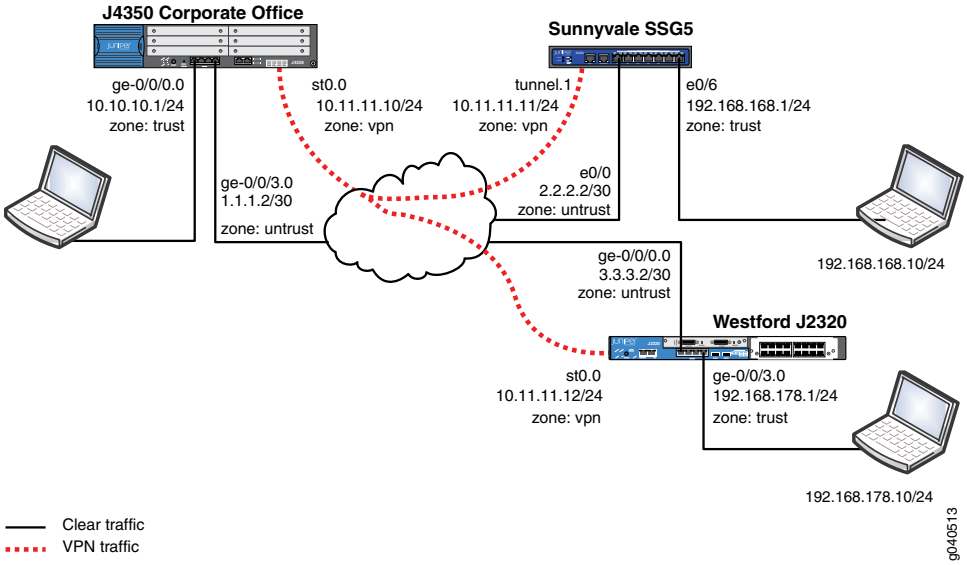

Required Settings

This example assumes the following settings:

- The internal LAN interface of the hub device (Corporate office) is **ge-0/0/0.0** in zone trust and has a private IP subnet.
- The Internet interface of the hub device (Corporate office) is **ge-0/0/3.0** in zone untrust and has a public IP subnet.
- The internal LAN interface of the spoke device (Westford office) is **ge-0/0/3.0** in zone trust and has a private IP subnet.
- The Internet interface of the spoke device (Westford office) is **ge-0/0/0.0** in zone untrust and has a public IP subnet.

NOTE: This example shows the configuration and verification of a multipoint interface in a hub-and spoke topology with two spokes. This example uses the following spokes as shown in Figure 2 on page 6:

- Spoke - device in Westford office, which is a Junos device running Junos OS Release 8.5 or later.
- Spoke - device in Sunnyvale office, which is an SSG device running ScreenOS 5.4.0 or later to outline interoperability requirements.

You can easily include additional spokes easily by duplicating the configurations from any existing spokes, changing IP addresses as needed, and adding any additional static routes for the new spoke local LANs.

- The secure tunnel interface st0.0 for the devices (in the Corporate office and in the Westford office) are configured in the **vpn** zone. This setting allows you to configure

Copyright © 2010, Juniper Networks, Inc.

unique policies specifically for tunnel (encrypted) traffic, while maintaining unique policies for clear (non-encrypted) traffic.

- All st0 interfaces of all peer devices have IP addresses configured within the same logical subnet. Configuring all peer tunnel interface IP addresses within the same logical subnet is recommended, but not mandatory. However, if you have configured OSPF with a point-to-multipoint link, then you must configure all peer tunnel interface IP addresses within the same logical subnet.
- Traffic is allowed in both directions from all remote offices (spokes) to the corporate LAN (hub). Traffic is also allowed from spoke to spoke. However, you can pass the traffic from one spoke to the other spoke only by first routing the traffic through the hub.
- A static NHTB entry is not required between the devices in the Westford office and the Corporate office because both devices operate on Junos OS. A static NHTB entry is required for the device in the Sunnyvale office, because the SSG device used here does not operate on Junos OS.
- The SSG5 has already been preconfigured with the correct settings.

Basic Configuration Steps for Hub and Spoke Devices

This topic includes the following sections:

Basic Steps to Configure the Hub (Device in Corporate Office)

Step-by-Step Procedure

The basic steps to configure the hub (device in the Corporate office) are:

1. Configure the IP addresses for the **ge-0/0/0.0**, **ge-0/0/3.0**, and **st0.0** interfaces.
2. Configure the default route to the Internet next hop and also configure static routes for each remote office LAN.

   NOTE: Optionally, you can use a dynamic routing protocol such as OSPF to configure the routes automatically, but that is beyond the scope of this example.

3. Configure the security zones, and bind the interfaces to the appropriate zones. Ensure that you have enabled the necessary host-inbound services on the interfaces or the zone. In this example, you must enable Internet Key Exchange (IKE) service on either the **ge-0/0/3** interface or to the zone untrust.
4. Configure address book entries for each zone.
5. Configure Phase 1 (IKE) gateway settings for both e remote offices. This example uses a standard proposal set. However, you can create a different proposal set, if required.

Copyright © 2010, Juniper Networks, Inc.

6. Configure Phase 2 IP Security- virtual private network (IPsec VPN) settings for both remote offices. Optionally, you can also configure the VPN monitor settings, if required. This example uses a standard proposal set, and Perfect Forward Secrecy (PFS) group 2. However, you can create a different proposal set, if required.
7. Bind the **st0.0** interface to the IPsec VPN.
8. Configure the **st0.0** for a multipoint interface. Configure the NHTB entries for any non-Junos spoke sites.

NOTE: If you are establishing a VPN between two Junos devices, then it is not necessary to configure an NHTB because the hub device can obtain the NHTB entry automatically during Phase 2 negotiations. However, if you have configured the VPN to establish tunnel on-traffic, then the hub site cannot initiate the VPN. Without an NHTB entry, the route for that remote peer would not be in active state.

In this scenario, either the tunnel must always be initiated from the spoke, or the hub device must have the establish-tunnel field configured to be activated immediately after the configuration changes are committed.

9. Configure security policies to permit remote office traffic into the host LAN (Corporate office) and vice versa.
10. Configure an outgoing **trust** to **untrust** permit-all policy with source Network Address Translation (NAT) for non-encrypted Internet traffic
11. Configure an intrazone policy in zone **vpn** to allow spokes to communicate with each other. Intrazone traffic is defined as traffic that ingresses (incoming) and egresses (outgoing) traffic through the same zone. By default intrazone traffic is denied.
12. Configure TCP maximum segment size (MSS) for IPsec traffic to eliminate the possibility of fragmented TCP traffic. This step reduces the resource usage on the device.

Basic Steps to Configure the Spoke (Device in Westford Office)

Step-by-Step Procedure

The basic steps to configure the spoke (device in the Westford office) are:

1. Configure the IP addresses for the **ge-0/0/0.0**, **ge-0/0/3.0** and **st0.0** interfaces.
2. Configure the default route to the Internet next hop, and also configure a static route for the Corporate office LAN.
3. Configure security zones and bind the interfaces to the appropriate zones. Also ensure that necessary host-inbound services are enabled on the interfaces or the zone.

   In this example, you must enable IKE service on either the **ge-0/0/0** interface or to the zone **untrust**.

Copyright © 2010, Juniper Networks, Inc.

4. Configure address book entries for each zone.
5. Configure Phase 1 (IKE) gateway settings. This example uses a standard proposal set.
6. Configure Phase 2 (IPsec) VPN settings.
7. Bind the **st0.0** interface to the IPsec VPN.

   This example uses a standard proposal set and PFS group 2.
8. Configure security policies to permit remote office (Westford office) traffic into the host LAN (Corporate office) and vice versa.
9. Configure an outgoing **trust** to **untrust** permit-all policy with source NAT for non-encrypted Internet traffic
10. Configure the TCP-MSS for IPsec traffic to eliminate the possibility of fragmented TCP traffic and to reduce the resource usage on the device

Example: Configuring the Multipoint VPN Configuration with Next-Hop Tunnel Binding

This topic includes the following sections:

Example: Configuring the Hub (Corporate Office)

Step-by-Step Procedure

1. Configure the IP addresses for the private LAN, public Internet, and the secure tunnel (**st0**) interfaces.

   NOTE: Junos OS uses the concept of units for the logical component of an interface. This example uses unit 0 and family inet (IPv4). We recommend configuring IP addresses for all peer devices within the same logical subnet (for st0 interfaces).

   ```
   [edit]
   user@hub# set interfaces ge-0/0/0 unit 0 family inet address 10.10.10.1/24
   user@hub# set interfaces ge-0/0/3 unit 0 family inet address 1.1.1.2/30
   user@hub# set interfaces st0 unit 0 family inet address 10.11.11.10/24
   ```

2. Configure a default route and other static routes for tunnel traffic.

   For a static route, you can normally specify the gateway IP address as the next-hop. For route-based VPNs with multipoint interface, you can specify the remote peer **st0** interface IP address as the next hop.

   ```
   [edit]
   user@hub# set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1

   user@hub# set routing-options static route 192.168.168.0/24 next-hop 10.11.11.11
   ```

Copyright © 2010, Juniper Networks, Inc.

```
user@hub# set routing-options static route 192.168.178.0/24 next-hop 10.11.11.12
```

3. Configure the security zones, and assign interfaces to the zones.

   Creating a unique zone for tunnel traffic enables you to create a specific set of policies for VPN traffic while maintaining a separate set of policies for non-VPN traffic. Also, you can create **deny** policies to prevent certain hosts from accessing the VPN.

   

   NOTE:

   No additional security policies are required, if:

   - You terminate the st0 interface in the same zone as the trusted LAN.
   - A policy is available that allows intrazone traffic in the same zone as the trusted LAN.

   ```
   [edit]
   user@hub# set security zones security-zone trust interfaces ge-0/0/0.0
   user@hub# set security zones security-zone untrust interfaces ge-0/0/3.0
   user@hub# set security zones security-zone vpn interfaces st0.0
   ```

4. Configure host-inbound services for each zone.

   Host-inbound services are for traffic destined for the Junos device. These settings include but are not limited to the FTP, HTTP, HTTPS, IKE, ping, rlogin, RSH, SNMP, SSH, Telnet, TFTP, and traceroute.

   In this example, we are assuming that all host-inbound services should be allowed from zone trust. For security reasons, we are allowing IKE only on the Internet-facing zone untrust, which is required for IKE negotiations to occur. However, you can enable other individual services, such as services for management, or services for troubleshooting, if required.

   ```
   [edit]
   user@hub# set security zones security-zone trust host-inbound-traffic
     system-services all
   user@hub# set security zones security-zone untrust host-inbound-traffic
     system-services ike
   ```

5. Configure the address book entries for each zone.

   In this example, we are using **addressbook** object names such as local-net, sunnyvale-net, and westford-net. Additional address book entries can added for any additional spokes, if required.

   ```
   [edit]
   user@hub# set security zones security-zone trust address-book address local-net
     10.10.10.0/24
   user@hub# set security zones security-zone vpn address-book address sunnyvale-net
     192.168.168.0/24
   user@hub# set security zones security-zone vpn address-book address westford-net
     192.168.178.0/24
   ```


Copyright © 2010, Juniper Networks, Inc.

6. Configure the IKE policy for main mode, predefined standard proposal set, and preshared key.

   In this example, we are using the standard proposal set, which includes the **esp-group2-3des-sha1** and **esp-group2- aes128-sha1** proposals. However, you may create a unique proposal and then specify it in the IKE policy in accordance with your corporate security policy.

   The same IKE policy can be used for all spoke VPNs, if needed.

   ```
   [edit]
   user@hub# set security ike policy ike-policy1 mode main
   user@hub# set security ike policy ike-policy1 proposal-set standard
   user@hub# set security ike policy ike-policy1 pre-shared-key ascii-text "secretkey"
   ```

7. Configure the spoke IKE gateways (Phase 1) with a peer IP address, an IKE policy, and an outgoing interface.

   A remote IKE peer can be identified by:

   - IP address
   - Fully qualified domain name / user-fully qualified domain name (FQDN/U-FQDN)
   - ASN1-DN (public key infrastructure [PKI] certificates)

   In this example, we are identifying the peer by IP address. Therefore the gateway address should be the remote peer's public IP address. You must also specify the correct external interface or peer ID to properly identify the IKE gateway during Phase 1 setup.

   ```
   [edit]
   user@hub# set security ike gateway westford-gate ike-policy ike-policy1
   user@hub# set security ike gateway westford-gate address 3.3.3.2
   user@hub# set security ike gateway westford-gate external-interface ge-0/0/3.0
   user@hub# set security ike gateway sunnyvale-gate ike-policy ike-policy1
   user@hub# set security ike gateway sunnyvale-gate address 2.2.2.2
   user@hub# set security ike gateway sunnyvale-gate external-interface ge-0/0/3.0
   ```

8. Configure the IPsec policy.

   In this example, we are using the standard proposal set, which includes the **esp-group2-3des-sha1** and **esp-group2- aes128-sha1** proposals. However, you may create a unique proposal and then specify it in the IPsec policy, if needed.

   ```
   [edit]
   user@hub# set security ipsec policy vpn-policy1 proposal-set standard
   user@hub# set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group2
   ```

9. Configure the IPsec VPN with an IKE gateway and IPsec policy, and bind them to the same **st0** interface.

   Binding an **st0** interface indicates that the VPN is a route-based VPN.

   You must specify an **st0** interface to successfully complete Phase 2 negotiations for route-based VPNs.

   ```
   [edit]
   user@hub# set security ipsec vpn westford-vpn ike gateway westford-gate
   ```

Copyright © 2010, Juniper Networks, Inc.

```
user@hub# set security ipsec vpn westford-vpn ike ipsec-policy vpn-policy1
user@hub# set security ipsec vpn westford-vpn bind-interface st0.0
user@hub# set security ipsec vpn sunnyvale-vpn ike gateway sunnyvale-gate
user@hub# set security ipsec vpn sunnyvale-vpn ike ipsec-policy vpn-policy1
user@hub# set security ipsec vpn sunnyvale-vpn bind-interface st0.0
```

10. Configure the **st0** interface as multipoint interface, then add a static NHTB entry for the spoke in the Sunnyvale office, which is an SSG device running ScreenOS. Because the spoke in the Sunnyvale office is not a Junos device, a static NHTB entry is required. You can also configure a static NHTB entry for another spoke in the Westford office if required (optional).

```
user@hub# set interfaces st0 unit 0 multipoint
user@hub# set interfaces st0 unit 0 family inet next-hop-tunnel 10.11.11.11 ipsec-vpn
  sunnyvale-vpn
```

11. Configure bidirectional security policies for tunnel traffic for all spokes.

    In this example, a security policy permits traffic in one direction but it also allows all reply traffic without the need for a reverse direction policy. However, since traffic may be initiated from either direction, bi-directional policies are required.

NOTE:

- If required, more granular policies can be created to permit/deny certain traffic.
- Because the policies are regular non-tunnel policies, they do not specify the IPsec profile.
- Source NAT can be enabled on the policy if desired, but that is beyond the scope of this example.
- If more spoke sites are added, you can add the additional source/destination match entries for the new spoke local LANs to permit the traffic.

```
[edit]
user@hub# edit security policies from-zone trust to-zone vpn

      [edit security policies from-zone trust to-zone VPN]
user@hub# set policy local-to-spokes match source-address local-net
user@hub# set policy local-to-spokes match destination-address sunnyvale-net
user@hub# set policy local-to-spokes match destination-address westford-net
user@hub# set policy local-to-spokes match application any
user@hub# set policy local-to-spokes then permit
Exit

[edit]
user@hub# edit security policies from-zone vpn to-zone trust

      [edit security policies from-zone vpn to-zone trust]
user@hub# set policy spokes-to-local match source-address sunnyvale-net
user@hub# set policy spokes-to-local match source-address westford-net
```


Copyright © 2010, Juniper Networks, Inc.

```
user@hub# set policy spokes-to-local match destination-address local-net
user@hub# set policy spokes-to-local match application any
user@hub# set policy spokes-to-local then permit
```

12. Configure a security policy for Internet traffic.

    A security policy is required to permit all traffic from zone trust to zone untrust.

    The device uses the specified **source-nat** interface, and translates the source IP address and port for outgoing traffic, using the IP address of the egress interface as the source IP address and using a random higher port for the source port. If required, you can create more granular policies to permit or deny certain traffic.

```
[edit]
user@hub# edit security policies from-zone trust to-zone untrust

      [edit security policies from-zone trust to-zone untrust]
user@hub# set policy any-permit match source-address any
user@hub# set policy any-permit match destination-address any
user@hub# set policy any-permit match application any
user@hub# set policy any-permit then permit source-nat interface
```

13. Configure an intrazone policy in the **vpn** zone for spoke-to-spoke traffic.

    This policy permits all traffic from zone **vpn** to zone **vpn**, which is intrazone traffic. You must configure an intra-zone policy to allow traffic through one spoke to another without dropping any traffic. If required, you can create more granular policies to permit or deny certain IP prefixes or protocols.

```
[edit]
user@hub# edit security policies from-zone vpn to-zone vpn

      [edit security policies from-zone vpn to-zone vpn]
user@hub# set policy spoke-to-spoke match source-address any
user@hub# set policy spoke-to-spoke match destination-address any
user@hub# set policy spoke-to-spoke match application any
user@hub# set policy spoke-to-spoke then permit
```

14. Configure the TCP MSS to eliminate the fragmentation of TCP traffic across the tunnel.

    TCP MSS is negotiated as part of the TCP three-way handshake. It limits the maximum size of a TCP segment to accommodate the maximum transmission unit (MTU) limits on a network. This is very important for VPN traffic, because the IPsec encapsulation overhead, along with the IP and Frame overhead, can cause the resulting ESP packet to exceed the MTU of the physical interface, resulting in fragmentation. Fragmentation increases the bandwidth and device resource usage, and should always be avoided. The recommended value for TCM MSS is 1350 for most Ethernet-based networks with an MTU of 1500 or higher. This value may need to be altered if any device in the path has a lower MTU value or if there is any added overhead such as PPP, or Frame Relay. As a general rule, you may need to experiment with different TCP MSS values to obtain optimal performance.

```
user@hub# set security flow tcp-mss ipsec-vpn mss 1350
```

Copyright © 2010, Juniper Networks, Inc.

Example: Configuring the Spoke (Westford Office)

Step-by-Step Procedure

1. Configure the IP addresses for the private LAN, public Internet, and the secure tunnel (**st0**) interfaces.

   NOTE: The steps for configuring the spoke device (Westford office) are similar to steps for configuring the hub device (Corporate office).

   NOTE: We recommend configuring IP addresses for all peer-devices within the same logical subnet (for st0 interfaces). Thus, the spoke device (Westford office) st0 interface is configured within the same subnet as the hub device (Corporate office) st0 interface.

   ```
   user@spoke# set interfaces ge-0/0/0 unit 0 family inet address 3.3.3.2/30
   user@spoke# set interfaces ge-0/0/3 unit 0 family inet address 192.168.178.1/24
   user@spoke# set interfaces st0 unit 0 family inet address 10.11.11.12/24
   ```

2. Configure a default route and other static routes for the tunnel traffic.

   Because the device in Westford office is in a spoke site, the **st0** interface type is point-to-point. Therefore, while configuring the next-hop option, you can specify the IP address of the **st0** interface on the hub site, or you can specify **st0.0** as the next hop.

   ```
   user@spoke# set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1
   user@spoke# set routing-options static route 10.10.10.0/24 next-hop 10.11.11.10
   user@spoke# set routing-options static route 192.168.168.0/24 next-hop 10.11.11.10
   ```

3. Configure the security zones and assign interfaces to the security zones.

   ```
   user@spoke# set security zones security-zone trust interfaces ge-0/0/3.0
   user@spoke# set security zones security-zone untrust interfaces ge-0/0/0.0
   user@spoke# set security zones security-zone vpn interfaces st0.0
   ```

4. Configure the host-inbound services for each zone.

   Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office).

   ```
   user@spoke# set security zones security-zone trust host-inbound-traffic
     system-services all

   user@spoke# set security zones security-zone untrust host-inbound-traffic
     system-services ike
   ```

5. Configure the address book entries for each zone.

   In this example we are using **address-book** object names such as local-net, sunnyvale-net, and westford-net.

Copyright © 2010, Juniper Networks, Inc.

You can add the additional spoke devices by:

- Creating an additional address book entry for each spokes' local LAN
- Creating a single address book entry that encompasses all spokes' local LANs

```
user@spoke# set security zones security-zone trust address-book address local-net
  192.168.178.0/24
user@spoke# set security zones security-zone vpn address-book address corp-net
  10.10.10.0/24
user@spoke# set security zones security-zone vpn address-book address
  sunnyvale-net 192.168.168.0/24
```

6. Configure the IKE policy for main mode, predefined standard proposal set, and preshared key.

NOTE: Details (proposal set and preshared key) used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office).

```
user@spoke# set security ike policy ike-policy1 mode main
user@spoke# set security ike policy ike-policy1 proposal-set standard
user@spoke# set security ike policy ike-policy1 pre-shared-key ascii-text "secretkey"
```

7. Configure the spoke IKE gateways (Phase 1) with a peer IP address, an IKE policy, and an outgoing interface.

NOTE:
Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office), except as follows:

- The external interface for the spoke device is ge-0/0/0.0.
- The peer address in the spoke device is the IP address of the hub device.

```
user@spoke# set security ike gateway corp-gate address 1.1.1.2
user@spoke# set security ike gateway corp-gate ike-policy ike-policy1
user@spoke# set security ike gateway corp-gate external-interface ge-0/0/0.0
```

8. Configure the IPsec policy for the standard proposal set.

NOTE: Details used in this step are the same for the both the spoke device (Westford office) and the hub device (Corporate office).

```
user@spoke# set security ipsec policy vpn-policy1 proposal-set standard
```

Copyright © 2010, Juniper Networks, Inc.

```
user@spoke# set security ipsec policy vpn-policy1 perfect-forward-secrecy keys
  group2
```

9. Configure the IPsec VPN with an IKE gateway and IPsec policy, and bind them to the same **st0** interface.

NOTE: Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office).

```
user@spoke# set security ipsec vpn corp-vpn ike gateway corp-gate
user@spoke# set security ipsec vpn corp-vpn ike ipsec-policy vpn-policy1
user@spoke# set security ipsec vpn corp-vpn bind-interface st0.0
```

10. Configure bidirectional security policies for tunnel traffic.

NOTE: Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office), except that the remote subnets used in the hub device local LAN and in the other spoke device local LAN are different.

```
[Edit]
user@spoke# edit security policies from-zone trust to-zone vpn
edit security policies from-zone trust to-zone vpn
user@spoke# set policy to-corp match source-address local-net
user@spoke# set policy to-corp match destination-address corp-net
user@spoke# set policy to-corp match destination-address sunnyvale-net
user@spoke# set policy to-corp match application any
user@spoke# set policy to-corp then permit
[Edit]
user@spoke# edit security policies from-zone vpn to-zone trust
edit security policies from-zone vpn to-zone trust
user@spoke# set policy from-corp match source-address corp-net
user@spoke# set policy from-corp match source-address sunnyvale-net
user@spoke# set policy from-corp match destination-address local-net
user@spoke# set policy from-corp match application any
user@spoke# set policy from-corp then permit
```

11. Configure a security policy for Internet traffic.

NOTE: Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office).

```
user@spoke# edit security policies from-zone trust to-zone untrust


user@spoke# set policy any-permit match source-address any
user@spoke# set policy any-permit match destination-address any
user@spoke# set policy any-permit match application any
user@spoke# set policy any-permit then permit source-nat interface
```


Copyright © 2010, Juniper Networks, Inc.

12. Configure the TCP MSS to eliminate the fragmentation of TCP traffic across the tunnel.

NOTE: Details used in this step are the same for both the spoke device (Westford office) and the hub device (Corporate office).

```
user@CORPORATE# set security flow tcp-mss ipsec-vpn mss 1350
```

Example: SSG Device Sample Configuration (For Reference Only)

Step-by-Step Procedure

1. This step provides information on SSG device configuration. Because the focus of this example is on Junos OS configuration and troubleshooting, the SSG device configuration is explained briefly.

   To show the configuration settings in Figure 2 on page 6, a sample of the relevant configurations is provided for an SSG5 device, strictly for reference.

   However, the concepts for configuring policy-based VPNs for Juniper Networks Firewall/VPN products are available in the Concepts and Examples (C&E) guides.

   For more information, see the *Concepts & Examples ScreenOS Reference Guide* available at http://www.juniper.net/techpubs/software/screenos/ .

```
user@SSG5# set zone name "VPN"
user@SSG5# set interface ethernet0/6 zone "Trust"
user@SSG5# set interface ethernet0/0 zone "Untrust"
user@SSG5# set interface "tunnel.1" zone "VPN"
user@SSG5# set interface ethernet0/6 ip 192.168.168.1/24
user@SSG5# set interface ethernet0/6 route
user@SSG5# set interface ethernet0/0 ip 2.2.2.2/30
user@SSG5# set interface ethernet0/0 route
user@SSG5# set interface tunnel.1 ip 10.11.11.11/24
user@SSG5# set flow tcp-mss 1350
user@SSG5# set address "Trust" "sunnyvale-net" 192.168.168.0 255.255.255.0
user@SSG5# set address "VPN" "corp-net" 10.10.10.0 255.255.255.0
user@SSG5# set address "VPN" "westford-net" 192.168.178.0 255.255.255.0
user@SSG5# set ike gateway "corp-ike" address 1.1.1.2 Main outgoing-interface
  ethernet0/0 preshare "secretkey" sec-level standard
user@SSG5# set vpn "corp-vpn" gateway "corp-ike" replay tunnel idletime 0
  sec-level standard
user@SSG5# set vpn "corp-vpn" bind interface tunnel.1
user@SSG5# set policy id 1 from "Trust" to "Untrust" "ANY" "ANY" "ANY" nat src
  permit
user@SSG5# set policy id 2 from "Trust" to "VPN" "sunnyvale-net" "corp-net"
  "ANY" permit
user@SSG5# set policy id 2
user@SSG5# set dst-address "westford-net"
exit
user@SSG5# set policy id 3 from "VPN" to "Trust" "corp-net" "sunnyvale-net"
  "ANY" permit
user@SSG5# set policy id 3
user@SSG5# set src-address "westford-net"
exit
```

Copyright © 2010, Juniper Networks, Inc.

```
user@SSG5# set route 10.10.10.0/24 interface tunnel.1
user@SSG5# set route 192.168.178.0/24 interface tunnel.1
user@SSG5# set route 0.0.0.0/0 interface ethernet0/0 gateway 2.2.2.1
```

Verification

This topic includes the following sections:

Confirm VPN Status

Purpose Confirm VPN status by checking the status of any IKE Phase 1 security associations.

Action Use the following command on the hub device (in the Corporate office)

```
user@host> show security ike security-associations

Index Remote Address State Initiator cookie Responder cookie Mode
6 3.3.3.2 UP 94906ae2263bbd8e 1c35e4c3fc54d6d3 Main
7 2.2.2.2 UP 7e7a1c0367dfe73c f284221c656a5fbc Main
```

Meaning The output indicates that:

- The remote peers (spokes) have the following IP addresses:
  - **3.3.3.2** (spoke device at the Westford office)
  - **2.2.2.2** (spoke device at the Sunnyvale office)
- The state showing **UP** for both remote peers indicates the successful association of Phase 1 establishment.
- The remote peer IP address, IKE policy, and external interfaces are all correct.
- Incorrect output would indicate that:
  - The remote peer status as Down.
  - There are no IKE security associations.
  - Incorrect IKE policy parameters such as wrong Mode type (Aggressive or Main), preshared keys, or Phase 1 proposals (all must match on both peers).

Copyright © 2010, Juniper Networks, Inc.

For more information, see “Troubleshooting Hub-and-Spoke VPNs” on page 24.

- Incorrect external interface.

  The external interface is invalid for receiving the IKE packets. Check the configurations for PKI-related issues or check the kmd log for any other errors or run traceoptions to find the mismatch. For more information, see “Troubleshooting Hub-and-Spoke VPNs” on page 24.

Get Peer Device’s Individual Index Numbers

Purpose Get details on the individual index numbers of the remote peer devices (spokes).

The Index number value is unique for each IKE SA for every remote peer.

Action

```
user@corporate> show security ike security-associations index 6 detail

IKE peer 3.3.3.2, Index 6,
Role: Responder, State: UP
Initiator cookie: 94906ae2263bbd8e, Responder cookie: 1c35e4c3fc54d6d3
Exchange type: Main, Authentication method: Pre-shared-keys
Local: 1.1.1.2:500, Remote: 3.3.3.2:500
Lifetime: Expires in 3571 seconds
Algorithms:
Authentication : sha1
Encryption : 3des-cbc
Pseudo random function: hmac-sha1
Traffic statistics:
Input bytes : 1128
Output bytes : 988
Input packets: 6
Output packets: 5
Flags: Caller notification sent
IPsec security associations: 1 created, 0 deleted
Phase 2 negotiations in progress: 1
Negotiation type: Quick mode, Role: Responder, Message ID: 1350777248
Local: 1.1.1.2:500, Remote: 3.3.3.2:500
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Flags: Caller notification sent, Waiting for done
```

Meaning The output displays the details of the spoke (in the Westford office) SA, such as the index, role (initiator or responder), status, exchange type, authentication method, encryption algorithms, traffic statistics, Phase 2 negotiation status, and so on.

You can use the output data to:

- Determine the role of the remote peer (spoke) device. Troubleshooting is easier when the peer device has the responder role.
- Obtain details regarding the authentication and encryption algorithms used.
- Obtain the traffic statistics to verify the traffic flow in both directions.
- Obtainthe number of IPsec SAs created or in progress

Copyright © 2010, Juniper Networks, Inc.

View IPsec (Phase 2) Security Associations

Purpose When IKE Phase 1 is confirmed, view the IPsec (Phase 2) security associations.

Action

```
user@corporate> show security ipsec security-associations

total configured sa: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<16384 2.2.2.2 500 ESP:3des/sha1 b2fc36f8 3564/ unlim - 0
>16384 2.2.2.2 500 ESP:3des/sha1 5d73929e 3564/ unlim - 0
total configured sa: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<16385 3.3.3.2 500 ESP:3des/sha1 70f789c6 28756/unlim - 0
>16385 3.3.3.2 500 ESP:3des/sha1 80f4126d 28756/unlim - 0
```

Meaning The output indicates that:

- There is a configured IPsec SA pair available. The port number 500 indicates that a standard IKE port is used. Otherwise, it is Network Address Translation-Traversal (NAT-T), 4500, or a random high port.
- The security parameter index (SPI) is used for both directions. The lifetime or usage limits of the SA are expressed either in seconds or in kilobytes. In the output, **28756/unlim** for **3.3.3.2** (spoke in the Westford office) indicates that the Phase 2 lifetime is set to expire in 28756 seconds and that there is no specified lifetime size.

  NOTE: The Phase 2 lifetime differ from the Phase 1 lifetime, because Phase 2 is not dependent on Phase 1 after the VPN is up.

- The Mon column refers to VPN monitoring status. A hyphen (-) in the Mon column indicates that VPN monitoring is not enabled for this SA. If VPN monitoring is enabled, then this field shows U (up) or D (down).

  NOTE: For information on VPN monitoring, refer to the complete documentation for Junos OS available at http://www.juniper.net/techpubs/ .

- The virtual system (vsys) is zero, which is the default value.
- The ID number shows the unique index value for each IPsec SA.

Display IPsec Security Association Details

Purpose Display the individual IPsec SA details identified by the index number for a remote peer device (Westford office).

The index value is unique for each IPsec SA. You can view more details for a particular SA by specifying the index value.

Copyright © 2010, Juniper Networks, Inc.

Action

```
user@corporate> show security ipsec security-associations index 16385 detail

Virtual-system: Root
Local Gateway: 1.1.1.2, Remote Gateway: 3.3.3.2
Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
DF-bit: clear
Direction: inbound, SPI: 1895270854, AUX-SPI: 0
Hard lifetime: Expires in 28729 seconds
Lifesize Remaining: Unlimited
Soft lifetime: Expires in 28136 seconds
Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: enabled, Replay window size: 32
Direction: outbound, SPI: 2163479149, AUX-SPI: 0
Hard lifetime: Expires in 28729 seconds
Lifesize Remaining: Unlimited
Soft lifetime: Expires in 28136 seconds
Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: enabled, Replay window size: 32
```

Meaning The output displays the local identity and the remote identity.

Note that a proxy ID mismatch may cause Phase 2 completion to fail. If no IPsec SA is listed, then confirm that the Phase 2 proposals, including the proxy ID settings, are correct for both peers. In this example, for route-based VPNs, the default proxy ID is **local=0.0.0.0/0**, **remote=0.0.0.0/0**, **service=any**.

NOTE:

- You might have to specify unique proxy IDs for each IPsec SA if you use multiple route-based VPNs from the same peer server's IP address.
- You might have to manually enter the proxy ID to match if you are using applications from some third-party vendors.
- You must specify st0 interface binding; otherwise, Phase 2 might fail to complete.

NOTE: If IPsec fails to complete, then check the kmd log or set traceoptions. For more information, see "Troubleshooting Hub-and-Spoke VPNs" on page 24.

Confirm Next-Hop Tunnel Bindings

Purpose After Phase 2 is complete for all peers, the next step to ensure that the routing is working properly, is to confirm that the NHTB table is established correctly.

To show the NHTB table, run the following command:

Copyright © 2010, Juniper Networks, Inc.

Action

```
user@corporate> show security ipsec next-hop-tunnels
Next-hop gateway interface IPsec VPN name Flag

10.11.11.11 st0.0 sunnyvale-vpn Static
10.11.11.12 st0.0 westford-vpn Auto
```

Meaning As in the network topology diagram in Figure 2 on page 6, the next-hop gateways are the IP addresses for the **st0** interfaces of all remote spoke peers. The next hop should be associated with the correct IPsec VPN name. Without an NHTB entry, it is not possible for the hub device to differentiate which IPsec VPN is associated with which next hop.

The flag can have one of the following options:

- Static – The NHTB is manually configured in the **st0.0** interface configurations, which is required if the peer device is not running Junos OS.
- Auto – The NHTB is not configured, but the entry was automatically populated into the table during Phase 2 negotiations between two Junos devices.

In this example, the NHTB table is not available on any of the devices in the spoke sites. This is because, from the spoke point of view, the **st0** interface is still a point-to-point link with only one IPsec VPN binding. Thus, if you use the **show security ipsec next-hop-tunnels** command on any of the devices in the spoke site (Westford office), you will not obtain any output.

Confirm Static Routes for Remote Peer Local LANs

Purpose For the NHTB table to be used, the static route needs to refer to the peer-devices (spoke) **st0** interface IP address. You can confirm the route to the remote peer by using the **show route <dest-ip-prefix>** operational command.

Action

```
user@corporate> show route 192.168.168.10
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.168.0/24 *[Static/5] 00:08:33
> to 10.11.11.11 via st0.0

user@corporate> show route 192.168.178.10

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.178.0/24 *[Static/5] 00:04:04
> to 10.11.11.12 via st0.0
```

Meaning In the output, the next hop is the remote peer **st0** interfaces' IP addresses and both routes point to **st0.0** as the outgoing interface.

Check Statistics and Errors for an IPsec SA

Purpose Check statistics and errors for an IPsec SA.

For troubleshooting purpose, check the Encapsulating Security Payload/Authentication Header (ESP/AH) counters for any errors with a particular IPsec SA.

Copyright © 2010, Juniper Networks, Inc.

Action

```
user@corporate> show security ipsec statistics index 16385
ESP Statistics:
Encrypted bytes: 920
Decrypted bytes: 6208
Encrypted packets: 5
Decrypted packets: 87
AH Statistics:
Input bytes: 0
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Meaning An error value of zero in the output indicates a normal condition.

We recommend running this command multiple times to observe any packet loss issues across a VPN. Output from this command also includes the statistics for encrypted and decrypted packet counters, error counters, and so on. You must enable security flow traceoptions to investigate which ESP packets are experiencing errors and why. For more information, see “Troubleshooting Hub-and-Spoke VPNs” on page 24.

Test Traffic Flow Across the VPN

Purpose Test traffic flow across the VPN after IKE Phase 1, Phase 2, routes, and NHTB entries have completed successfully. You can test traffic flow by using the **ping** command. You can ping from local host to remote host. You can also initiate pings from the Junos device itself.

This example shows how to initiate a ping request from the Junos device to the remote host at the Sunnyvale office. You can use the same procedure to ping a host device at the Westford office to confirm connectivity. Note that when pings are initiated from the Junos device, the source interface must be specified to ensure that the correct route lookup takes place and that the appropriate zones are referenced in the policy lookup.

In this example, the **ge-0/0/0.0** interface resides in the same security zone as the local host and must be specified in the ping request so that the policy lookup can be from zone trust to zone untrust.

Action

```
user@CORPORATE> ping 192.168.168.10 interface ge-0/0/0 count 5
PING 192.168.168.10 (192.168.168.10): 56 data bytes
64 bytes from 192.168.168.10: icmp_seq=0 ttl=127 time=8.287 ms
64 bytes from 192.168.168.10: icmp_seq=1 ttl=127 time=4.119 ms
64 bytes from 192.168.168.10: icmp_seq=2 ttl=127 time=5.399 ms
64 bytes from 192.168.168.10: icmp_seq=3 ttl=127 time=4.361 ms
64 bytes from 192.168.168.10: icmp_seq=4 ttl=127 time=5.137 ms
--- 192.168.168.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.119/5.461/8.287/1.490 ms
```

Confirm the Connectivity

Purpose Confirm the connectivity between a remote host and a local host.

Copyright © 2010, Juniper Networks, Inc.

Action

```
ssg5-> ping 10.10.10.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 10.10.10.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/4/5 ms
ssg5-> ping 192.168.178.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.178.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=8/8/10 ms
```

Meaning You can confirm end-to-end connectivity from a remote host at a spoke site to a local host at the Corporate office LAN by using the **ping** command. In this example, the command is initiated from the SSG5 device.

Failed end-to-end connectivity may indicate an issue with routing, policy, end host, or encryption/decryption of the ESP packets. To verify the exact causes of the failure:

- Check the IPsec statistics for details on errors as described in “Check Statistics and Errors for an IPsec SA” on page 22.
- Confirm end host connectivity by using the ping command from a host on the same subnet as the end host. If the end host is reachable by other hosts, then you can assume that the issue is not with the end host.
- Enable security flow traceoptions for troubleshooting the routing-and -policy-related issues.

Troubleshooting Hub-and-Spoke VPNs

The basic troubleshooting steps are as follows:

- Identifying and isolating the problem
- Debugging the problem

The common approach to begin troubleshooting is to start with the lowest layer of the Open Systems Interconnection (OSI) layers and working your way up the OSI stack to determine in which layer the failure occurs. The steps for troubleshooting are as follows:

- Confirm the physical connectivity of the Internet link at the physical and data link levels.
- Confirm that the Junos device has connectivity to the Internet next hop and connectivity to the remote IKE peer.
- Confirm IKE Phase 1 completion.
- Confirm IKE Phase 2 completion if IKE Phase 1 completion is successful.
- Confirm the traffic flow across the VPN (if the VPN is up and active).

Junos OS includes the traceoptions feature. Using this feature, you can enable a traceoption flag to write the data from the trace to a log file. The log file may be predetermined, or manually configured and the file is stored in flash memory. These trace

 Copyright © 2010, Juniper Networks, Inc.

logs can be retained even after a system reboot. Check the available flash storage before implementing traceoptions.

You can enable the traceoptions feature in configuration mode and commit the configuration to use the traceoptions feature. Similarly, to disable traceoptions, you must deactivate traceoptions in configuration mode and commit the configuration.

If the VPN is not in the UP state then there is very little reason to test any transit traffic across the VPN. Likewise if Phase 1 is not successful, then there is no need to look at Phase 2 issues.

Check the Free Disk Space on Your Device

Problem You need to check the statistics on the free disk space in your device file systems to make sure that you have enough memory available to perform other tasks.

Solution Use **show system storage** command output to verify the free disk space.

```
user@corporate> show system storage
Filesystem Size Used Avail Capacity Mounted on
/dev/ad0s1a 213M 136M 75M 65% /
devfs 1.0K 1.0K 0B 100% /dev
devfs 1.0K 1.0K 0B 100% /dev/
/dev/md0 144M 144M 0B 100% /junos
/cf 213M 136M 75M 65% /junos/cf
devfs 1.0K 1.0K 0B 100% /junos/dev/
procfs 4.0K 4.0K 0B 100% /proc
/dev/bo0s1e 24M 13K 24M 0% /config
/dev/md1 168M 7.3M 147M 5% /mfs
/dev/md2 58M 38K 53M 0% /jail/tmp
/dev/md3 7.7M 108K 7.0M 1% /jail/var
devfs 1.0K 1.0K 0B 100% /jail/dev
/dev/md4 1.9M 6.0K 1.7M 0% /jail/html/oem
```

The **/dev/ad0s1a** represents the onboard flash memory and is currently at 65% capacity.

NOTE: You can view the available system storage in the J-Web interface under the System Storage option.

NOTE: You can enable traceoptions to log the trace data to the filenames specified or to the default log file to receive the output of the tracing operation. The output of the traceoptions is placed in /var/log/kmd.

Check the Log Files to Verify Different Scenarios and to Upload Log Files to an FTP Server

Problem You need to the check the log files to verify that logging to the syslog is working and that there are no errors shown in the security IKE debug messages and security flow debug messages.

Copyright © 2010, Juniper Networks, Inc.

Solution To verify the messages in the syslog, use the **show log kmd**, **show log security-trace**, and **show log messages** commands.

```
user@corporate> show log kmd
user@corporate> show log security-trace
user@corporate> show log messages
```

NOTE: You can view a list of all logs in the /var/log directory by using the show log command.

Log files can also be uploaded to an FTP server by using the **file copy** command.

(operational mode):

```
user@corporate> file copy /var/log/kmd ftp://10.10.10.10/kmd.log
ftp://10.10.10.10/kmd.log 100% of 35 kB 12 MBps
```

Enable IKE Traceoptions to View Messages on IKE

Problem You need to view additional details Phase 1 and Phase 2 negotiation issues and error messages for by enabling IKE and PKI traceoptions.

To verify success or failure messages for IKE or IPsec, you can view the key management process (kmd) log by using the **show log kmd** command. Because the kmd log displays some general messages, it may be useful to obtain additional details by enabling IKE and PKI traceoptions.

NOTE: Generally, it is best practice to troubleshoot the peer that has the responder role. You must obtain the trace output from the initiator and the responder to understand the cause of a failure.

Solution You can enable IKE traceoptions by configuring the file to write trace options and setting the flag for trace messages in **edit security ike traceoptions** hierarchy.

```
user@corporate> configure
Entering configuration mode
[edit]
user@corporate> edit security ike traceoptions
[edit security ike traceoptions]
```

```
user@corporate# set file ?
Possible completions:
<filename> Name of file in which to write trace information
files Maximum number of trace files (2..1000)
match Regular expression for lines to be logged
no-world-readable Don't allow any user to read the log file
size Maximum trace file size (10240..1073741824)
world-readable Allow any user to read the log file

[edit]
```


Copyright © 2010, Juniper Networks, Inc.

```
[edit security ike traceoptions]
user@corporate# set flag ?

Possible completions:
all Trace everything
certificates Trace certificate events
database Trace security associations database events
general Trace general events
ike Trace IKE module processing
parse Trace configuration processing
policy-manager Trace policy manager processing
routing-socket Trace routing socket messages
timer Trace internal timer events
```

NOTE: If you do not specify file names for the <filename> field, then all IKE traceoptions are written to the kmd log.

To write trace data to the log, you must specify at least one flag option. For example:

- file size — Maximum size of each trace file, in bytes. For example 1m or 1000000 can generate a maximum file size of 1 MB.
- file files — Maximum number of trace files to be generated and stored in flash memory.

NOTE: To start the trace, you must commit your configuration.

Setting Up IKE Traceoptions to Troubleshoot IKE-Related Issues

Problem You need to configure the recommended settings for IKE traceoptions, such as file size, policy-manager, flag, and routing-socket.

Solution You can configure the required IKE traceoptions in the **edit security ike traceoptions** hierarchy by using the following commands:

- **set file size 1m**
- **set flag policy-manager**
- **set flag ike**
- **set flag routing-socket**

```
user@corporate# edit security ike traceoptions
[edit security ike traceoptions]
user@corporate# set file size 1m
user@corporate# set flag policy-manager
user@corporate# set flag ike
user@corporate# set flag routing-socket
user@corporate# commit
```

Copyright © 2010, Juniper Networks, Inc.

Analyzing the Phase 1 and Phase 2 Success Messages

Problem Confirm the success of Phase 1 and Phase 2.

Solution Use the **show log kmd** command output to confirm that IKE Phase 1 and Phase 2 conditions are successful, as shown below:

```
Oct 8 10:41:40 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=1.1.1.2)
remote=ipv4(udp:500,[0..3]=2.2.2.2)
Oct 8 10:41:51 Phase-2 [responder] done for p1_local=ipv4(udp:500,[0..3]=1.1.1.2)
p1_remote=ipv4(udp:500,[0..3]=2.2.2.2)
p2_local=ipv4_subnet(any:0,[0..7]=10.10.10.0/24)
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
```

The sample output indicates:

- **1.1.1.2** — Local address.
- **2.2.2.2** — Remote address.
- **udp: 500** — Indicates that no NAT-T was negotiated.
- **Phase 1 [responder] done** — Indicates the Phase 1 status, along with the role (initiator or responder).
- **Phase 2 [responder] done** — Indicates the Phase 1 status with proxy ID information.

You can also confirm the IPsec SA status by using the verification commands mentioned in “Display IPsec Security Association Details” on page 20.

Analyzing the Phase 1 Failure Message (Proposal Mismatch)

Problem Phase 1 (responder) fails with an error because of proposal mismatch.

Solution To resolve this issue, ensure that the parameters for the Phase 1 proposals match on both the responder and the initiator.

The following sample shows output from the **show log kmd** command, where the IKE Phase 1 condition is a failure caused by mismatched proposal:

```
Oct 8 10:31:10 Phase-1 [responder] failed with error(No proposal chosen) for
local=unknown(any:0,[0..0]=) remote=ipv4(any:0,[0..3]=2.2.2.2)

Oct 8 10:31:10 1.1.1.2:500 (Responder) <-> 2.2.2.2:500 { 011359c9 ddef501d - 2216ed2a
 bfc50f5f [-
1] / 0x00000000 } IP; Error = No proposal chosen (14)
```

The sample output indicates:

- **1.1.1.2** — Local address.
- **2.2.2.2** — Remote address.
- **udp: 500** — No NAT-T was negotiated.
- **Phase-1 [responder] failed with error (No proposal chosen)** — Phase 1 failure caused by proposal mismatch.

Copyright © 2010, Juniper Networks, Inc.

Analyzing the Phase 1 Failure Message (Policy Lookup Failure)

Problem Phase 1 (responder) fails with error caused by a policy lookup failure.

This condition indicates that Phase 1 is failing because of a proposal mismatch and because the responder is not recognizing the incoming request as originating from a valid gateway peer. The peer was not recognized because of an incorrect peer address, a mismatched peer ID type, or an incorrect peer ID, depending on whether this is a dynamic VPN or a static VPN.

Solution To resolve this issue, configure the following before Phase 1:

- Correct peer IP address on the local peer
- Local peer with an IKE ID type of IP address

The following sample shows the output from the **show log kmd** command, where the Phase 1 failure is caused by a policy lookup failure:

```
Oct 8 10:39:40 Unable to find phase-1 policy as remote peer:2.2.2.2 is not recognized .

Oct 8 10:39:40 KMD_PM_P1_POLICY_LOOKUP_FAILURE: Policy lookup for Phase-1
[responder] failed for
p1_local=ipv4(any:0,[0..3]=1.1.1.2) p1_remote=ipv4(any:0,[0..3]=2.2.2.2)
Oct 8 10:39:40 1.1.1.2:500 (Responder) <-> 2.2.2.2:500 { 18983055 dbe1d0af -
a4d6d829 f9ed3bba [-
1] / 0x00000000 } IP; Error = No proposal chosen (14)
```

The sample output indicates:

- **1.1.1.2** — Local address.
- **2.2.2.2** — Remote address.
- **Unable to find phase-1 policy as remote peer:2.2.2.2 is not recognized** — This indicates a Phase 1 failure caused by a proposal mismatch and by the responder's not recognizing the incoming request as originating from a valid gateway peer (peer:2.2.2.2 is not recognized).

Analyzing the Phase 1 Failure Message (Invalid Payload Type)

Problem Phase 1 (responder) fails because of an invalid payload type. The invalid payload type indicates a problem with IKE packet decryption caused by mismatch of the preshared keys.

Solution To resolve this issue, configure the preshared keys to match on the peers.

The following sample shows the output from the **show log kmd** command, when the Phase 1 condition is a failure caused by invalid payload type:

```
Oct 8 10:36:20 1.1.1.2:500 (Responder) <-> 2.2.2.2:500 { e9211eb9 b59d543c - 766a826d
 bd1d5ca1 [-
1] / 0x00000000 } IP; Invalid next payload type = 17
```

Copyright © 2010, Juniper Networks, Inc.

```
Oct 8 10:36:20 phase-1 [responder] failed with error(Invalid payload type) for
local=unknown(any:0,[0..0]=) remote=ipv4(any:0,[0..3]=2.2.2.2)
```

The sample output indicates:

- 1.1.1.2 — Local address.
- 2.2.2.2 — Remote address.
- Phase 1 [responder] failed with error (invalid payload type) — Indicates Phase 1 failure caused by invalid payload type.

Analyzing the Phase 2 Failure Message (Proposal Mismatch)

Problem Phase 2 fails with error caused by proposal mismatch between two peers.

Solution To resolve this issue, configure the Phase 2 proposals to match on the peers.

The following sample shows output of the **show log kmd** command, when the Phase 2 condition is a failure:

```
Oct 8 10:53:34 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=1.1.1.2)
remote=ipv4(udp:500,[0..3]=2.2.2.2)
Oct 8 10:53:34 1.1.1.2:500 (Responder) <-> 2.2.2.2:500 { cd9dff36 4888d398 -
6b0d3933 f0bc8e26 [0]
/ 0x1747248b } QM; Error = No proposal chosen (14)
```

The sample output indicates:

- **1.1.1.2** — Local address.
- **2.2.2.2** — Remote address.
- **Phase 1 [responder] done** — Indicates Phase 1 success.
- **Error = No proposal chosen** — Indicates that no proposal was chosen during Phase 2 negotiations.

Analyzing the Phase 2 Failure Message (Proxy ID Mismatch)

Problem Phase 2 fails with error caused by proxy ID mismatch between two peers, resulting from a mismatch of configurations on the local peer.

Solution To resolve this issue, configure the proxy ID on one of the peers so that it matches the other peer.

The following sample shows output from the **show log kmd** command, when the Phase 2 condition is a failure.

```
Oct 8 10:56:00 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=1.1.1.2)
remote=ipv4(udp:500,[0..3]=2.2.2.2)
Oct 8 10:56:00 Failed to match the peer proxy ids
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
p2_local=ipv4_subnet(any:0,[0..7]=10.10.20.0/24)
 for the remote peer:ipv4(udp:500,[0..3]=2.2.2.2)
Oct 8 10:56:00 KMD_PM_P2_POLICY_LOOKUP_FAILURE: Policy lookup for Phase-2 [responder]
failed for
```


Copyright © 2010, Juniper Networks, Inc.

```
p1_local=ipv4(udp:500,[0..3]=1.1.1.2) p1_remote=ipv4(udp:500,[0..3]=2.2.2.2)
p2_local=ipv4_subnet(any:0,[0..7]=10.10.20.0/24)
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
Oct 8 10:56:00 1.1.1.2:500 (Responder) <-> 2.2.2.2:500 { 41f638eb cc22bbfe -
43fd0e85 b4f619d5 [0]
/ 0xc77fafcf } QM; Error = No proposal chosen (14)
```

The sample output indicates:

- **1.1.1.2** — Local address.
- **2.2.2.2** — Remote address.
- **Phase 1 [responder] done** — Indicates Phase 1 success.
- **Policy lookup for Phase 2 [responder] failed** — Indicates that the incorrect proxy IDs are received. In the sample output, the two proxy IDs received are **192.168.168.0/24** (remote) and **10.10.20.0/24** (local) (for service=any). Based on the configuration example in "Example: Configuring the Spoke (Westford Office)" on page 14, the expected local address is 10.10.10.0/24. This shows that there is a mismatch of configurations on the local peer, resulting in the failure of proxy ID match.

NOTE: Note that for a route-based VPN, the proxy ID by default is all zeroes (local=0.0.0.0/0, remote=0.0.0.0/0, service=any). If the remote peer specifies a proxy ID other than all zeroes, then you must manually configure the proxy ID within the IPsec profile of the peer.

Common Problems Related to IKE and PKI

Problem Troubleshoot common problems related to IKE and PKI.

Solution Enabling the traceoptions feature helps you to gather more information for debugging issues than is obtainable from the normal log entries. You can use the traceoptions log to understand the reasons for traffic not passing through the tunnel because of problems related to route lookup, security policy, or some other flow issue (assuming that the IPsec tunnel is up).

Details of flow traceoption output are beyond the scope of this example. For more information on traceoptions output, see the *Junos Enhanced Services Route-Based VPN Configuration and Troubleshooting* at http://kb.juniper.net/kb/documents/public/junos_es/Junos_ES_Route_based_VPN_to_ScreenOS.pdf.

NOTE: Enabling the flow traceoptions increases the device resource usage, and it should always be best avoided during peak traffic load times or when CPU usage is very high.

We recommend enabling packet filters to reduce the resource usage and to facilitate classification of the required packets.

Copyright © 2010, Juniper Networks, Inc.

NOTE: We recommend deleting or deactivating all flow traceoptions and removing any unnecessary log files from flash memory after completing troubleshooting. To disable traceoptions, you must deactivate traceoptions in configuration mode and then commit the configuration.

Enable Flow Traceoption to View Messages on Routing or Policy Issues

Problem View the log messages on routing- or policy-related issues.

Solution You need to the check the log files to verify that logging to the syslog is working and that there are no errors in the security IKE debug messages or the security flow debug messages.

```
user@corporate# edit security flow traceoptions
[edit security flow traceoptions]

user@corporate# set file ?
Possible completions:
<filename> Name of file in which to write trace information
files Maximum number of trace files (2..1000)
match Regular expression for lines to be logged
no-world-readable Don't allow any user to read the log file
size Maximum trace file size (10240..1073741824)
world-readable Allow any user to read the log file

[edit security flow traceoptions]

user@corporate# set flag ?
Possible completions:
ager Ager events
all All events
basic-datapath Basic packet flow
cli CLI configuration and commands changes
errors Flow errors
fragmentation Ip fragmentation and reassembly events
high-availability Flow high-availability information
host-traffic Flow host-traffic information
lookup Flow lookup events
multicast Multicast flow information
packet-drops Packet drops
route Route information
session Session creation and deletion events
session-scan Session scan information
tcp-advanced Advanced TCP packet flow
tcp-basic TCP packet flow
tunnel Tunnel information
```

NOTE: If you do not specify filenames for the filename field, then all flow traceoptions are written to the security-trace log. However, you can specify a different filename, if desired.

Copyright © 2010, Juniper Networks, Inc.

To write data to the log, you must specify at least one flag option. For example:

- file size — Maximum size of each trace file, in bytes. For example, 1m or 1000000 can generate a maximum file size of 1 MB.
- files — Maximum number of log files to be generated and stored in flash memory.

NOTE: To start the trace, you must first commit the configuration.

Configure Packet Filters to Reduce the Resource Usage

Problem Limit the scope of the traffic to be captured by the flow traceoptions.

Solution You can limit the scope of traffic captured by configuring packet filters as shown in the following command.

[edit security flow traceoptions]

```
user@corporate# set packet-filter filter-name ?
Possible completions:
+ apply-groups Groups from which to inherit configuration data
+ apply-groups-except Don't inherit configuration data from these groups
destination-port Match TCP/UDP destination port
destination-prefix Destination IPv4 address prefix
interface Logical interface
protocol Match IP protocol type
source-port Match TCP/UDP source port
source-prefix Source IPv4 address prefix
```

NOTE: Note the following points concerning the packet-filter:

- By configuring the packet-filter option, you can filter the output based on source or destination IP, source or destination port, interface, and IP protocol.
- You can configure up to 64 filters.
- A packet filter can work in reverse direction to capture the reply traffic, if the source of the original packet matches the filter. For more information on flow packet-filter options, see “Troubleshoot the Traffic, Using Traceoptions with Packet Filters” on page 33. Terms listed within the same packet filter act as a Boolean logical AND statement. That means that all statements within the packet filter need to match in order to write the output to the log. A listing of multiple filter names acts as a logical OR.

Troubleshoot the Traffic, Using Traceoptions with Packet Filters

Problem Troubleshoot the traffic flow from the remote peer (Westford Office) to the local host (Corporate Office).

Copyright © 2010, Juniper Networks, Inc.

Solution You can troubleshoot the traffic flow between the remote peer to the local host by using packet filters. You can use the output details from each flow traceoption command (as shown in Table 2 on page 34) to analyze the traffic.

```
user@corporate# edit security flow traceoptions
[edit security flow traceoptions]
user@corporate# set file size 1m files 3
user@corporate# set flag basic-datapath
user@corporate# set packet-filter remote-to-local source-prefix 192.168.178.0/24
user@corporate# set packet-filter remote-to-local destination-prefix 10.10.10.0/24
user@corporate# set packet-filter local-to-remote source-prefix 10.10.10.0/24
user@corporate# set packet-filter local-to-remote destination-prefix 192.168.178.0/24
user@corporate# set packet-filter remote-esp protocol 50
user@corporate# set packet-filter remote-esp source-prefix 3.3.3.2/32
user@corporate# commit
```

Table 2 on page 34 provides output details for each flow traceoption setting in this sample configuration.

Table 2: Output Details for Flow Traceoption Setting

Output of the Settings	What it indicates...
[edit security flow traceoptions] user@CORPORATE# show file flow-trace-log size 1m files 3; flag basic-datapath;	• Because there are no filename is specified, all flow traceoptions are written to the **security-trace** log file. • The **security-trace** log file is set to 1 MB and up to 3 files can be created. Because the flow traceoption may generate a large log file to capture the traffic, a single file may not be adequate. • Flag **basic-datapath** captures the details for most flow-related problems.

packet-filter
remote-to-local {
source-prefix
192.168.168.0/24;
destination-prefix
10.10.10.0/24;
}
	• The **packet-filter remote-to-local** is set for capturing the decapsulated or unencrypted traffic from the remote peer to the local host.
	• The filter acts as a Boolean logical AND statement because there are multiple terms listed. This filter captures the packets only if the source IP address and destination IP address match. If one of the addresses does not match, then the packet is not captured.
	• The packet filters are bidirectional, and it is not necessary to configure a filter for the reply traffic.

packet-filter
local-to-remote {
source-prefix
10.10.10.0/24;
destination-prefix
192.168.178.0/24;
}
	The **packet-filter local-to-remote** is required even though it is not required to set a filter for capturing the reply traffic. However a filter can capture only the packets that are originally sourced from the specified side. Thus, the **local-to-remote** filter in this example may still be required to capture traffic from the local side to the remote side.

Copyright © 2010, Juniper Networks, Inc.

Table 2: Output Details for Flow Traceoption Setting *(continued)*

Output of the Settings	What it indicates...

packet-filter remote-esp {
protocol 50;
source-prefix 3.3.3.2/32;
}
	• The **packet-filter remote-esp** is optional and depends on whether or not the previous filter could capture any packets.
	• This filter can capture all ESP (IP protocol 50) or encrypted packets from remote peer 2.2.2.2.
	NOTE: Because this filter is configured at the bottom of the hierarchy, it captures all encrypted traffic from server 2.2.2.2, which may not be required.
	NOTE: However, the last filter can capture all encrypted traffic from 2.2.2.2 including packets are not required. Since the last filter captures unencrypted traffic, this filter may not be required.

Thus, using the filters, you can troubleshoot any traffic flow issues to and from the Corporate Office and the Westford site. Additional filters can be configured for troubleshooting from Westford to Sunnyvale and vice versa. In addition, to help narrow the scope, a single host can be specified with the /32 mask to avoid having too much data written to the trace log. Finally, as always, if any assistance is needed in interpreting the data from any of the traceoption logs, contact your regional JTAC (Juniper Technical Assistance Center). The JTAC website can be found at: http://www.juniper.net/customers/support/ .

Related Documentation

- Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview on page 1
- Verifying Hub-and-Spoke VPN Configuration on page 37

Copyright © 2010, Juniper Networks, Inc.

 Copyright © 2010, Juniper Networks, Inc.

Verifying Hub-and-Spoke VPN Configuration

This topic includes the following sections:

Verifying Configuration of the Hub (Device in Corporate Office)

Use the **show configuration** command to verify the configuration.

```
system {
    host-name CORPORATE;
    root-authentication {
        encrypted-password "$1$0wc5IQiB$MTQlktoQ9/nRF10Gntin./"; ## SECRET-DATA
    }
    services {
        ssh;
        web-management {
            http {
                interface ge-0/0/0.0;
            }
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any any;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 10.10.10.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 1.1.1.2/30;
            }
        }
    }
    st0 {
        unit 0 {
```

Copyright © 2010, Juniper Networks, Inc.

```
        multipoint;
        family inet {
          next-hop-tunnel 10.11.11.11 ipsec-vpn sunnyvale-vpn;
          address 10.11.11.10/24;
        }
      }
    }
}
routing-options {
    static {
      route 0.0.0.0/0 next-hop 1.1.1.1;
      route 192.168.168.0/24 next-hop 10.11.11.11;
      route 192.168.178.0/24 next-hop 10.11.11.12;
    }
}
security {
    ike {
      traceoptions {
        flag policy-manager;
        flag ike;
        flag routing-socket;
        flag general;
      }
      policy ike-policy1 {
        mode main;
        proposal-set standard;
        pre-shared-key ascii-text "$9$LrN7w2mPQF/t24jqmfn6rev"; ## SECRET-DATA
      }
      gateway sunnyvale-gate {
        ike-policy ike-policy1;
        address 2.2.2.2;
        external-interface ge-0/0/3.0;
      }
      gateway westford-gate {
        ike-policy ike-policy1;
        address 3.3.3.2;
        external-interface ge-0/0/3.0;
      }
    }
    ipsec {
      policy vpn-policy1 {
        perfect-forward-secrecy {
          keys group2;
        }
        proposal-set standard;
      }
      vpn sunnyvale-vpn {
        bind-interface st0.0;
        ike {
          gateway sunnyvale-gate;
          ipsec-policy vpn-policy1;
        }
      }
      vpn westford-vpn {
        bind-interface st0.0;
        ike {
```

 Copyright © 2010, Juniper Networks, Inc.

```
        gateway westford-gate;
        ipsec-policy vpn-policy1;
      }
    }
  }
  zones {
    security-zone trust {
      address-book {
        address local-net 10.10.10.0/24;
      }
      host-inbound-traffic {
        system-services {
          all;
        }
      }
      interfaces {
        ge-0/0/0.0;
      }
    }
    security-zone untrust {
      host-inbound-traffic {
        system-services {
          ike;
        }
      }
      interfaces {
        ge-0/0/3.0;
      }
    }
    security-zone vpn {
      address-book {
        address sunnyvale-net 192.168.168.0/24;
        address westford-net 192.168.178.0/24;
      }
      interfaces {
        st0.0;
      }
    }
  }
  policies {
    from-zone trust to-zone untrust {
      policy any-permit {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit {
            source-nat {
              interface;
            }
          }
        }
      }
    }
```

Copyright © 2010, Juniper Networks, Inc.

```
    from-zone trust to-zone vpn {
      policy local-to-spokes {
        match {
          source-address local-net;
          destination-address [ sunnyvale-net westford-net ];
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone vpn to-zone trust {
      policy spokes-to-local {
        match {
          source-address [ sunnyvale-net westford-net ];
          destination-address local-net;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone vpn to-zone vpn {
      policy spoke-to-spoke {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
  }
  flow {
    tcp-mss {
      ipsec-vpn {
        mss 1350;
      }
    }
  }
}
```

NOTE: In the preceding sample of output from the **show configuration** command, the highlighted lines show traceoption configurations for troubleshooting purposes.

TIP: Delete or deactivate the traceoptions after you complete troubleshooting.

Copyright © 2010, Juniper Networks, Inc.

Verifying Configuration of the Spoke (Device in Westford Office)

Use the **show configuration** command to verify the configuration.

```
system {
    host-name Westford;
    root-authentication {
        encrypted-password "$1$Qk3dVh9X$d3KOf3dhR6uQKhi8FWU.P0"; ## SECRET-DATA
    }
    services {
        web-management {
            http {
                interface ge-0/0/0.0;
            }
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any any;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 3.3.3.2/30;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 192.168.178.1/24;
            }
        }
    }
    st0 {
        unit 0 {
            family inet {
                address 10.11.11.12/24;
            }
        }
    }
}
routing-options {
    static {
        route 0.0.0.0/0 next-hop 1.1.1.1;
```

Copyright © 2010, Juniper Networks, Inc.

```
        route 10.10.10.0/24 next-hop 10.11.11.10;
        route 192.168.168.0/24 next-hop 10.11.11.10;
    }
}
security {
    ike {
        traceoptions {
            flag policy-manager;
            flag ike;
            flag routing-socket;
            flag general;
        }
        policy ike-policy1 {
            mode main;
            proposal-set standard;
            pre-shared-key ascii-text "$9$VNsaGF39A0IGDPQFnpu8X7"; ## SECRET-DATA
        }
        gateway corp-gate {
            ike-policy ike-policy1;
            address 1.1.1.2;
            external-interface ge-0/0/0.0;
        }
    }
    ipsec {
        policy vpn-policy1 {
            perfect-forward-secrecy {
                keys group2;
            }
            proposal-set standard;
        }
        vpn corp-vpn {
            bind-interface st0.0;
            ike {
                gateway corp-gate;
                ipsec-policy vpn-policy1;
            }
        }
    }
    zones {
        security-zone trust {
            address-book {
                address local-net 192.168.178.0/24;
            }
            host-inbound-traffic {
                system-services {
                    all;
                }
            }
            interfaces {
                ge-0/0/3.0 {
                }
            }
        }
        security-zone untrust {
            host-inbound-traffic {
                system-services {
```


Copyright © 2010, Juniper Networks, Inc.

```
            ike;
          }
        }
        interfaces {
          ge-0/0/0.0 {
          }
        }
      }
      security-zone vpn {
        address-book {
          address corp-net 10.10.10.0/24;
          address sunnyvale-net 192.168.168.0/24;
        }
        interfaces {
          st0.0;
        }
      }
    }
    policies {
      from-zone trust to-zone untrust {
        policy any-permit {
          match {
            source-address any;
            destination-address any;
            application any;
          }
          then {
            permit {
              source-nat {
                interface;
              }
            }
          }
        }
      }
      from-zone vpn to-zone trust {
        policy from-corp {
          match {
            source-address [ corp-net sunnyvale-net ];
            destination-address local-net;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone trust to-zone vpn {
        policy to-corp {
          match {
            source-address local-net;
            destination-address [ corp-net sunnyvale-net ];
            application any;
          }
          then {
            permit;
```

Copyright © 2010, Juniper Networks, Inc.

```
            }
          }
        }
      }
      flow {
        tcp-mss {
          ipsec-vpn {
            mss 1350;
          }
        }
      }
    }
```

NOTE: In the preceding sample of output from the show configuration command, the highlighted lines show traceoption configurations for troubleshooting purposes.

TIP: Delete or deactivate the traceoptions after you complete troubleshooting.

Related Documentation

- Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview on page 1
- Example: Configuring Hub-and-Spoke VPNs using Next-Hop Tunnel Binding on page 5

Copyright © 2010, Juniper Networks, Inc.