



Point-to-Point GRE over IPsec Design Guide

OL-9023-01

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Networking Academy logo, Cisco Unity, Fast Step, Follow Me Browsing, FormShare, FrameShare, IGX, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ Logo, iQ Net Readiness Scorecard, MGX, the Networkers logo, ScriptBuilder, ScriptShare, SMARTnet, TransPath, Voice LAN, Wavelength Router, and WebViewer are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and Discover All That's Possible are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, GigaStack, IOS, IP/TV, LightStream, MICA, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0110R)

Point-to-Point GRE over IPsec Design Guide

© 2006 Cisco Systems, Inc. All rights reserved.



Preface 1-7

- Introduction 1-7
 - Target Audience 1-8
 - Scope of Work 1-8
- Document Organization 1-9

CHAPTER 1

Point-to-Point GRE over IPsec Design Overview 1-1

- Starting Assumptions 1-1
- Quality of Service per p2p GRE Tunnel Interface 1-2
- Design Components 1-2
- Topology 1-2
- Headend System Architectures 1-3
 - Single Tier Headend Architecture 1-3
 - Dual Tier Headend Architecture 1-4
 - Single Tier Headend Architecture versus Dual Tier Headend Architecture 1-4
- Branch Router Considerations 1-7
 - Static p2p GRE over IPsec with a Branch Static Public IP Address 1-7
 - Static p2p GRE over IPsec with a Branch Dynamic Public IP Address 1-7
- High Availability 1-7
- Best Practices and Known Limitations 1-7
 - Best Practices Summary 1-8
 - Known Limitations Summary 1-9

CHAPTER 2

Point-to-Point GRE over IPsec Design and Implementation 2-1

- Design Considerations 2-1
 - Topology 2-2
 - Headend System Architectures 2-2
 - Single Tier Headend Architecture 2-3
 - Dual Tier Headend Architecture 2-4
 - IP Addressing 2-5
 - Generic Route Encapsulation 2-6
 - GRE Keepalives 2-6
 - Using a Routing Protocol across the VPN 2-7
 - Route Propagation Strategy 2-7

- Crypto Considerations 2-7
 - IPsec Tunnel versus Transport Mode 2-8
 - Dead Peer Detection 2-8
- Configuration and Implementation 2-8
 - ISAKMP Policy Configuration 2-8
 - Dead Peer Detection Configuration 2-9
 - IPsec Transform and Protocol Configuration 2-10
 - Access Control List Configuration for Encryption 2-11
 - Crypto Map Configuration 2-12
 - Applying Crypto Maps 2-13
 - Tunnel Interface Configuration—Branch Static Public IP Address 2-14
 - Tunnel Interface Configuration—Branch Dynamic Public IP Address 2-14
 - GRE Keepalive Configuration 2-15
 - Routing Protocol Configuration 2-16
 - Route Propagation Configuration 2-17
- High Availability 2-17
 - Common Elements in all HA Headend Designs 2-18
 - 1+1 (Active-Standby) Failover Headend Resiliency Design 2-18
 - Load Sharing with Failover Headend Resiliency Design 2-21
 - N+1 Failover Architecture 2-22
 - Dual Tier Headend Architecture Effect on Failover 2-23
- QoS 2-23
- IP Multicast 2-23
- Interactions with Other Networking Functions 2-23
 - Network Address Translation and Port Address Translation 2-24
 - Dynamic Host Configuration Protocol 2-24
 - Firewall Considerations 2-24
 - Headend or Branch 2-24
 - Firewall Feature Set and Inbound ACL 2-25
 - Double ACL Check Behavior (Before 12.3(8)T) 2-25
 - Crypto Access Check on Clear-Text Packets Feature (12.3(8)T and Later) 2-25
- Common Configuration Mistakes 2-26
 - Crypto Peer Address Matching using PSK 2-26
 - Transform Set Matches 2-26
 - ISAKMP Policy Matching 2-26

CHAPTER 3

Scalability Considerations 3-1

- General Scalability Considerations 3-1
- IPsec Encryption Throughput 3-1

Packets Per Second—Most Important Factor	3-2
Tunnel Quantity Affects Throughput	3-2
GRE Encapsulation Affects Throughput	3-2
Routing Protocols Affect CPU Overhead	3-2
Headend Scalability	3-3
Tunnel Aggregation Scalability	3-3
Aggregation Scalability	3-4
Customer Requirement Aggregation Scalability Case Studies	3-4
Customer Example with 300–500 Branches	3-4
Customer Example with 1000 Branches	3-5
Customer Example with 1000–5000 Branches	3-8
Branch Office Scalability	3-9

CHAPTER 4

Scalability Test Results (Unicast Only)	4-1
Scalability Test Bed Network Diagram	4-1
Scalability Test Methodology	4-3
Headend Scalability Test Results—p2p GRE over IPsec	4-3
Headend Scalability Test Results—p2p GRE Only	4-4
Branch Office Scalability Test Results	4-4
AES versus 3DES Scalability Test Results	4-5
Failover and Convergence Performance	4-6
Software Releases Evaluated	4-7

CHAPTER 5

Case Studies	5-1
Static p2p GRE over IPsec with a Branch Dynamic Public IP Address Case Study	5-1
Overview	5-1
Sample Topology	5-2
Addressing and Naming Conventions	5-2
Configuration Examples	5-4
p2p GRE Tunnel and Interface Addressing	5-4
Crypto Map Configurations (Crypto Tunnel)	5-5
Headend EIGRP Configuration	5-6
Verification	5-6
Summary	5-7
Moose Widgets Case Study	5-7
Customer Overview	5-7
Design Considerations	5-9
Preliminary Design Considerations	5-9
Sizing the Headend	5-10

Sizing the Branch Sites 5-10
Tunnel Aggregation and Load Distribution 5-11
Network Layout 5-11

APPENDIX A

Scalability Test Bed Configuration Files A-1

Cisco 7200VXR Headend Configuration A-1
Cisco Catalyst 6500/Sup2/VPNSM Headend Configuration A-2
Cisco 7600/Sup720/VPN SPA Headend Configuration (p2p GRE on Sup720) A-6
Cisco 7600/Sup720/VPN SPA Headend Configuration (p2p GRE on VPN SPA) A-10
Cisco 7200VXR/7600 Dual Tier Headend Architecture Configurations A-13
Cisco 7600/Sup720/VPN SPA Headend Configuration A-17
ISR Branch Configuration A-19

APPENDIX B

Legacy Platform Test Results B-1

Cisco Headend VPN Routers (Legacy) B-1
Cisco Branch Office VPN Routers (Legacy) B-1

APPENDIX C

References and Reading C-1

APPENDIX D

Acronyms D-1



Preface

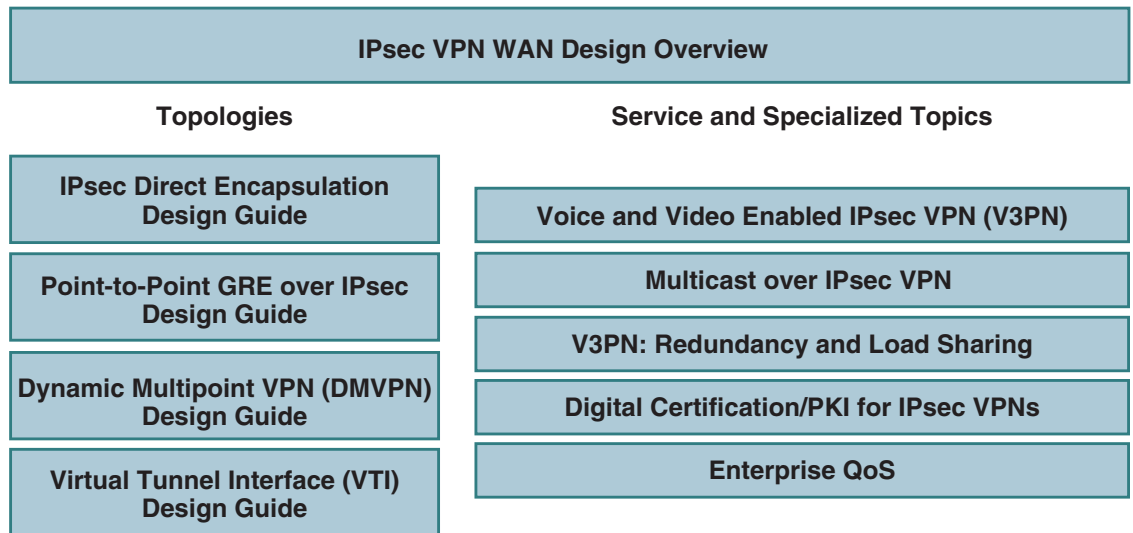
This design guide defines the comprehensive functional components required to build a site-to-site virtual private network (VPN) system in the context of enterprise wide area network (WAN) connectivity. This design guide covers the design topology of point-to-point (p2p) Generic Route Encapsulation (GRE) over IP Security (IPsec).

This design guide is part of an ongoing series that addresses VPN solutions, using the latest VPN technologies from Cisco, and based on practical design principles that have been tested to scale.

Introduction

Figure 1 lists the IPsec VPN WAN architecture documentation.

Figure 1 IPsec VPN WAN Architecture Documentation



190897

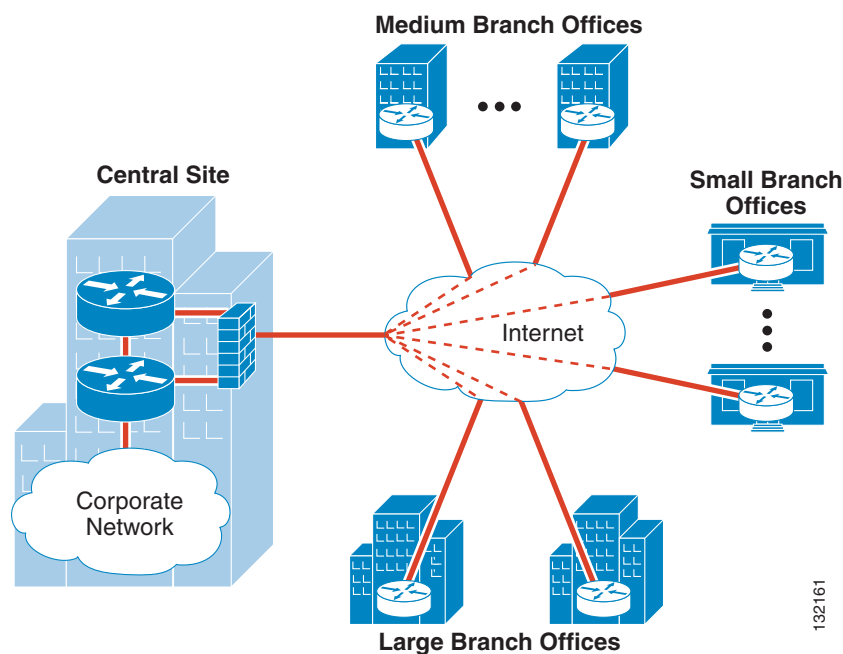
The IPsec VPN WAN architecture is divided into multiple design guides based on technologies. These guides are available at the following URL: <http://www.cisco.com/go/srnd>.

Each technology uses IPsec as the underlying transport mechanism for each VPN. The operation of IPsec is outlined in the *IPsec VPN WAN Design Overview*. The reader must have a basic understanding of IPsec before reading further. The *IPsec VPN WAN Design Overview* also outlines the criteria for selecting a specific IPsec VPN WAN technology. This document should be used to select the correct technology for the proposed network design.

This document serves as a design guide for those intending to deploy a site-to-site VPN based on IPsec and GRE. This version of the design guide focuses on Cisco IOS VPN router products.

The primary topology discussed is a hub-and-spoke design, where the primary enterprise resources are located in a large central site, with a number of smaller sites or branch offices connected directly to the central site over a VPN. A high-level diagram of this topology is shown in [Figure 2](#).

Figure 2 Hub-and-Spoke VPN



This design guide begins with an overview, followed by design recommendations, as well as product selection and performance information. Finally, a case study and configuration examples are presented.

Target Audience

This design guide is targeted for systems engineers and provides guidelines and best practices for customer deployments.

Scope of Work

This version of the design guide addresses the following applications of the solution:

- Cisco VPN routers running IOS
- p2p GRE tunneling over IPsec is the tunneling method
- Site-to-site VPN topologies

- Use of Enhanced Interior Gateway Routing Protocol (EIGRP) as a routing protocol across the VPN with GRE configurations
- Dynamic crypto peer address with static GRE endpoints
- Dead Peer Detection (DPD)
- Converged data and voice over IP (VoIP) traffic requirements
- Quality of service (QoS) features are enabled
- Evaluation of Cisco VPN product performance in scalable and resilient designs

Document Organization

This guide contains the chapters in the following table.

Section	Description
Chapter 1, “Point-to-Point GRE over IPsec Design Overview.”	Provides an overview of the VPN site-to-site design topology and characteristics.
Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation.”	Provides an overview of some general design considerations that need to be factored into the design, followed by sections on implementation, high availability, QoS, and IP multicast.
Chapter 3, “Scalability Considerations.”	Provides guidance in selecting Cisco products for a VPN solution, including sizing the headend, choosing Cisco products that can be deployed for headend devices, and product sizing and selection information for branch devices.
Chapter 4, “Scalability Test Results (Unicast Only).”	Provides test results from the Cisco test lab to provide design guidance on the scalability of various platforms in p2p GRE over IPsec VPN configurations.
Chapter 5, “Case Studies.”	Provides two case studies as reference material for implementing p2p GRE over IPsec designs.
Appendix A “Scalability Test Bed Configuration Files.”	Provides the configurations for the central and branch sites.
Appendix B “Legacy Platform Test Results.”	Provides scalability test results for legacy products.
Appendix C “References and Reading.”	Provides references to further documentation.
Appendix D “Acronyms.”	Provides definitions for acronyms.



Point-to-Point GRE over IPsec Design Overview

This chapter provides an overview of the VPN site-to-site design topology and characteristics. [Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation,”](#) provides more detail on the design considerations. [Chapter 3, “Scalability Considerations,”](#) presents Cisco product options for deploying the design.

Starting Assumptions

The design approach presented in this design guide makes the following starting assumptions:

- The design supports a typical converged traffic profile for customers (see [Chapter 4, “Scalability Test Results \(Unicast Only\).”](#))
- It is assumed that the customer has a need for diverse traffic requirements, such as IP multicast, multiprotocol, and support for routing. The use of p2p GRE and a routing protocol are also discussed in more detail in [Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation.”](#)
- Cisco products should be maintained at reasonable CPU utilization levels. This is discussed in more detail in [Chapter 3, “Scalability Considerations,”](#) including recommendations for both headend and branch routers, and software revisions.
- Although costs were certainly considered, the design recommendations assume that the customer will deploy current VPN technologies, including hardware-accelerated encryption.
- Voice over IP (VoIP) and video are assumed to be requirements in the network. Detailed design considerations for handling VoIP and other latency-sensitive traffic are not explicitly addressed in this design guide, but may be found in *Voice and Video Enabled IPsec VPN (V3PN)*, which is available at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns241/c649/ccmigration_09186a00801ea79c.pdf
- Finally, this design is targeted for deployment by enterprise-owned VPNs; however, the concepts and conclusions are valid regardless of the ownership of the edge tunneling equipment, and are therefore valuable for service provider-managed VPNs as well.

Quality of Service per p2p GRE Tunnel Interface

All headend scalability results were performed on various Cisco platforms. Results were obtained without a hierarchical class-based weighted fair queueing service policy applied per p2p GRE tunnel interface. This QoS configuration is a parent service policy that shapes packets in the logical tunnel interface and then queues packets within the shaped rate in a child service policy.

Applying a hierarchical class-based weighted fair queueing service policy to the p2p GRE tunnel interface may require packets to be process-switched while the shaper is active; thus, it is CPU intensive.

**Note**

The *Virtual Tunnel Interface (VTI) Design Guide* and the *Dynamic Multipoint VPN (DMVPN) Design Guide* (<http://www.cisco.com/go/srnd>), provide performance guidance on per tunnel interface QoS.

Design Components

VPNs have many applications, including extending reachability of an enterprise WAN, or replacing classic WAN technologies such as leased lines, Frame Relay, and ATM. Site-to-site VPNs are primarily deployed to connect branch office locations to the central site (or sites) of an enterprise.

The requirements of enterprise customers for traditional private WAN services such as multiprotocol support, high availability, scalability, and security, are also requirements for VPNs. VPNs can often meet these requirements more cost-effectively and with greater flexibility than private WAN services.

The key components of this site-to-site VPN design are the following:

- Cisco high-end VPN routers serving as VPN headend termination devices at a central campus (headend devices)
- Cisco VPN access routers serving as VPN branch termination devices at branch office locations (branch devices)
- p2p GRE over IPsec to perform headend-to-branch interconnections
- Internet services procured from a third-party ISP (or ISPs) serving as the WAN interconnection medium

Cisco VPN routers are a good choice for site-to-site VPN deployments because they can accommodate any network requirement inherited from a Frame Relay or private line network, such as support for IP multicast and latency-sensitive traffic, routing for resiliency, and support for non-IP protocols such as IPX or SNA. See [Chapter 3, “Scalability Considerations,”](#) for a discussion on selection of headend and branch products.

Topology

In a p2p GRE over IPsec design, the following three topologies can be implemented:

- Hub-and-spoke
- Partial mesh
- Full mesh

The hub-and-spoke topology is discussed in this design guide because it is the most widely deployed.


Headend System Architectures

This section describes the two headend system architectures that can be implemented, depending on the scalability requirements.

Single Tier Headend Architecture

In a Single Tier Headend Architecture, both the p2p GRE and crypto functionally co-exist on the same router CPU. [Figure 1-1](#) shows this hub-and-spoke topology.

Figure 1-1 Single Tier Headend Architecture



[Figure 1-1](#) shows a hub-and-spoke network with multiple headend devices for redundancy. Headends are p2p GRE and crypto tunnel aggregation routers servicing multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, headends can advertise branch routes using IP routing protocols such as EIGRP or OSPF.

Dual Tier Headend Architecture

In a Dual Tier Headend Architecture, the p2p GRE and crypto do not functionally co-exist on the same router CPU. Figure 1-2 shows this hub-and-spoke topology.

Figure 1-2 Dual Tier Headend Architecture




Figure 1-2 shows a hub-and-spoke network with multiple headend devices for redundancy. p2p GRE headends as well as crypto headends together service multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, the p2p GRE headends can advertise branch routes using IP routing protocols such as EIGRP or OSPF.

Single Tier Headend Architecture versus Dual Tier Headend Architecture

The choice between the Single Tier Headend Architecture and the Dual Tier Headend Architecture depends on the following criteria:

- Topology
- Performance
- Value

In either architecture, the topology being designed is the first consideration. For this comparison, a hub-and-spoke topology is the only topology detailed because this discussion is focused on the headend router. Table 1-1 lists the technical limitations of the two architectures.

Table 1-1 Single Tier Headend versus Dual Tier Headend Architecture— Technical Limitations

Headend Architecture	Router	Crypto Configuration	Crypto IP Address	GRE Configuration	GRE IP Address	Tunnel Protection
Single Tier	Headend	Static or dynamic	Static	p2p GRE static	Static	Optional

Table 1-1 Single Tier Headend versus Dual Tier Headend Architecture— Technical Limitations

	Branch	Static	Static or dynamic	p2p GRE static	Static	Optional
Dual Tier	Headend	Static or dynamic	Static	p2p GRE static	Static	Not valid
	Branch	Static	Static or dynamic	p2p GRE static	Static	Not valid

Tunnel protection requires the same source and destination IP address for both the GRE tunnel and the crypto tunnel; therefore, it is not a valid option in a Dual Tier Headend Architecture. Both architectures support all of the options above as a hub-and-spoke topology.

When considering performance and value, the two should be considered together. Performance is based on the number of packets a router can forward in a given timeframe, or packets per second (pps). Value is the price for a specific router based on the pps rate. Given these two considerations, [Table 1-2](#) shows both price and performance for the primary headend router choices currently available.

Table 1-2 Platform Price and Performance

Platform	pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1	p2p GRE only—200 Kpps	\$20,000
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	\$30,000
Cisco 7606 with Sup720, SSC400, and VPN SPA	p2p GRE and 3DES—520 Kpps	\$100,000
Cisco 7606 with Sup720, SSC400, and VPN SPA	3DES only—600 Kpps	\$100,000

For the purpose of this design guide, *these prices are not the actual price* of the specific platforms because prices change on a periodic basis. The prices are provided as a reference to show the value of one architecture over the other. The reader should obtain specific pricing per platform when comparing the platforms. However, the pps values have been verified in Cisco testing.

With this in mind, now consider the value for each of the architectures, given the performance numbers stated for a p2p GRE over IPsec design. The Single Tier Headend Architecture is the easiest to demonstrate because both the p2p GRE and encryption processes are housed on a single routing processor. [Table 1-3](#) shows these results.

Table 1-3 Single Tier Headend Architecture Comparison

Platform	pps (bi-directional)	Quantity required	Aggregate pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	3	120 Kpps (40 Kpps * 3)	\$90,000 (\$30,000 * 3)

Table 1-3 Single Tier Headend Architecture Comparison

Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	13	520 Kpps (40 Kpps * 13)	\$390,000 (\$30,000 * 13)
Cisco 7606 with Sup720, SSC400, and VPN SPA	p2p GRE and 3DES—520 Kpps	1	520 Kpps	\$100,000

Note from this comparison that the break-even point regarding value per performance is approximately 120 Kpps at \$100,000. Below this price point, the clear choice is to deploy multiple Cisco 7200VXRs with NPE-G1 and Dual SA-VAM2+; however, above this point, the decision is much different. To obtain the same 520 Kpps with the Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+ solution, it requires thirteen chassis at a cost of \$390,000, versus a single Cisco 7606 with Sup720, SSC400, and VPN SPA at \$100,000. The difference represents a substantial capital savings over the long term. Also, support contracts are increased from one to thirteen as well.

The Dual Tier Headend Architecture is slightly harder to demonstrate because the p2p GRE and encryption processes are housed on separate routing processors. [Table 1-4](#) shows these results.

Table 1-4 Dual Tier Headend Architecture Comparison

Platform	pps (bi-directional)	Quantity required	Aggregate pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	5	200 Kpps (40 Kpps * 5)	\$150,000 (\$30,000 * 5)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	15	600 Kpps (40 Kpps * 15)	\$450,000 (\$30,000 * 15)
Cisco 7606 with Sup720, SSC400, and VPN SPA for encryption only and Cisco 7200VXR with NPE-G1 for p2p GRE <i>only</i>	p2p GRE and 3DES—200 Kpps	1 Cisco 7606 and 1 Cisco 7200VXR	200 Kpps (200 Kpps p2p GRE only on the Cisco 7200VXR is the limitation)	\$120,000 (\$100,000 Cisco 7606 with Sup720, SSC400, and VPN SPA) + (\$20,000 Cisco 7200VXR with NPE-G1)
Cisco 7606 with Sup720, SSC400, and VPN SPA for encryption only and Cisco 7200VXR with NPE-G1 for p2p GRE <i>only</i>	p2p GRE and 3DES—600 Kpps	1 Cisco 7606 and 3 Cisco 7200VXRs	600 Kpps (200 Kpps p2p GRE on each Cisco 7200VXR with 600 Kpps on a single VPN SPA)	\$160,000 (\$100,000 Cisco 7606 with Sup720, SSC400, and VPN SPA) + (3 * \$20,000 Cisco 7200VXR with NPE-G1)

Note from this comparison that the break-even point regarding value per performance is approximately 200 Kpps at \$150,000. Below this price point, the clear choice is to deploy multiple Cisco 7200VXRs with NPE-G1 and Dual SA-VAM2+; however, above this point, the decision is much different. To obtain the same 600 Kpps with the Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+ solution, it requires fifteen chassis at a cost of \$450,000, versus a single Cisco 7606 with Sup720, SSC400, and VPN SPA with three Cisco 7200VXRs with NPE-G1 at \$160,000. The difference represents a substantial capital savings over the long term. Also, support contracts are increased from 4 to 15 as well.

Branch Router Considerations

Branches are typically access routers that provide p2p GRE over IPsec tunnel(s) from the branch office locations to the central site. In addition to terminating the VPN tunnels, the branch router often provides WAN access, and in some implementations may serve as a firewall.

Static p2p GRE over IPsec with a Branch Static Public IP Address

In this scenario, the public IP address of the branch router is a statically defined IP address. Both the p2p GRE and crypto tunnels are sourced from this statically defined public IP address.

Static p2p GRE over IPsec with a Branch Dynamic Public IP Address

In this scenario, the public IP address of the branch router is a dynamically assigned IP address. The p2p GRE tunnel is sourced from a loopback interface with an administratively assigned private IP address. The crypto tunnel is sourced from the dynamically assigned public IP address. A static host route is required in the headend router to ensure p2p GRE packets destined to the branch router loopback interface are encrypted.

High Availability

Network resiliency is provided differently depending on the initial network requirements. This design guide uses a dynamic IGP routing protocol across the VPN. Because IPsec does not provide the ability to run protocols requiring IP multicast (such as EIGRP), it is necessary to use p2p GRE in conjunction with IPsec. p2p GRE supports more diverse traffic across the VPN, including IP multicast and non-IP protocols.

For high availability in the case of a failure, each branch access router should have two p2p GRE over IPsec tunnels, a primary and secondary, provisioned to different headend tunnel aggregation routers. This is discussed further in [High Availability, page 2-17](#).

Best Practices and Known Limitations

The following sections contain a summary of the best practices and limitations for the design. More detailed information is provided in [Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation.”](#)

Best Practices Summary

The following list summarizes the best practices for a p2p GRE over IPsec design, supporting multiprotocol and/or IP multicast traffic including routing protocols:

- General best practices
 - Use IPsec in tunnel mode for best flexibility.
 - Configure Triple DES (3DES) or AES for encryption of transported data (exports of encryption algorithms to certain countries may be prohibited by law).
 - Implement Dead Peer Detection (DPD) to detect loss of communication between peers.
 - Deploy hardware-acceleration of IPsec to minimize router CPU overhead, to support traffic with low-latency/jitter requirements, and for the highest performance for cost.
 - Keep IPsec packet fragmentation to a minimum on the customer network by setting MTU size or using PMTU Discovery (PMTUD).
 - Use Digital Certificates/PKI for scalable tunnel authentication keys.
 - Set up QoS service policies as appropriate on headend and branch router interfaces to ensure performance of latency-sensitive applications (for more information, see the *V3PN QoS and IPsec Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.)
 - Configure a routing protocol such as EIGRP or OSPF with route summarization for dynamic routing.
- Headend best practices
 - Configure dynamic crypto maps on the headend to support dynamically addressed branches and to simplify provision of new branches.
 - If high availability is a requirement, implement a design with redundancy of headend equipment and WAN circuits.
 - Distribute branch office tunnels across a number of headend routers to balance loading and aggregation capacity of the hub(s).
 - Select Cisco VPN router products at the headend based on considerations for the following:
 - Number of tunnels to be aggregated
 - Maximum throughput in both pps and bps to be aggregated
 - Performance margin for resiliency and failover scenarios
 - Maintaining CPU utilization below design target
 - See [Chapter 3, “Scalability Considerations,”](#) for more information.
- Branch office best practices
 - Configure multiple p2p GRE over IPsec tunnels to redundant headends.
 - Select Cisco VPN router products at the branch offices based on considerations for the following:
 - Maximum throughput in both pps and bps
 - Allowances for other integrated services that may be running on the router, such as firewall, IPS, and NAT/PAT
 - Maintaining CPU utilization below 65–80 percent
 - See [Chapter 3, “Scalability Considerations,”](#) for more information.

- The QoS pre-classify feature is desirable in VPN designs where both QoS and IPsec occur on the same system. The network manager should verify correct operation.

Known Limitations Summary

The following lists at a high level the known limitations for a p2p GRE over IPsec VPN design:

- General limitations
 - p2p GRE acceleration is currently limited to 2047 tunnels with the VPN SPA and VPNSM, and 2000 tunnels with the Sup720. Other factors such as the number of sustainable routing peers may affect the maximum number of tunnels in a design.
 - Although IPsec can typically scale to thousands of tunnels on some platforms, a routed p2p GRE over IPsec design is generally limited by the routing protocol being used and the number of routing peers exchanging routing information, such as the following:
 - 500 for the Cisco 7200VXR with NPE-G1
 - 600 for the Cisco 7200VXR with NPE-G2
 - 1000 for the Cisco 7600 (or Catalyst 6500) with Sup720
 - There are significant scalability limitations for supporting IP multicast over p2p GRE over IPsec designs. For more information, see the *Multicast over IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.
 - QoS pre-classify is not a supported feature in the Cisco Catalyst 6500 or on the Cisco 7600 platforms.
 - p2p GRE over IPsec designs do not readily support a “touchless” headend configuration concept, and therefore generally require configuration changes to the headend router(s) to provision new branch office tunnels.
 - GRE keepalives are not currently functional on the Cisco 7600 because of DDTS.
- Single Tier Headend Architecture limitations
 - It is possible to implement a QoS service policy at the tunnel/destination level to achieve per-VPN tunnel QoS and prevent hub-to-spoke overruns. However, the number of branches that can be supported in this configuration is limited if shaping is engaged concurrently for a large percentage of the configured tunnels.
- Dual Tier Headend Architecture limitations
 - Tunnel protection is not supported.
- Branch office limitations
 - The p2p GRE over IPsec tunnel must be initiated by the remote branch in cases where remote routers acquire their address via a dynamically served IP address. The crypto headend cannot initiate the tunnel to the branch. However, if a dynamic routing protocol is configured, the hello packets initiated by the branch router force the tunnel to be established when the branch router is started, and the tunnel is maintained by this control plane traffic.
 - In designs with QoS and IPsec, interaction between QoS and IPsec anti-replay can result in dropped packets if packets delayed by QoS fall outside the anti-replay sequence number window at the receiver.

Additional detailed information on these recommendations is discussed in the chapters that follow.



Point-to-Point GRE over IPsec Design and Implementation

In designing a VPN deployment for a customer, it is essential to integrate broader design considerations such as high availability, resiliency, IP multicast, and quality of service (QoS).

This chapter starts with an overview of some general design considerations that need to be factored into the design, followed by sections on implementation, high availability, QoS, and IP multicast.

Design Considerations

Headend sites are typically connected with DS3, OC3, or even OC12 bandwidth, while branch offices may be connected by fractional T1, T1, T3, or increasingly, broadband DSL or cable access.

To provide redundancy, the branch router should have two or more tunnels to the campus headends. These headend routers can be geographically separated or co-located. For maximum protection, both headend and site redundancy should be implemented. This design guide focuses on a solution with only two point-to-point (p2p) GRE tunnels per branch terminating to two headend routers, to simplify the routing domain.

The IPsec control plane uses dynamic crypto maps at the headend to minimize configuration changes in the event of new branches being added. Dynamic crypto maps are also implemented to support branches with a dynamic Internet address as their crypto peer. Dead Peer Detection (DPD) is configured to perform automatic detection of ISAKMP peer loss, thus tearing down the VPN tunnel. Alternatively, the IPsec tunnel protection feature can be configured on tunnel interfaces.

The GRE tunnel uses p2p GRE on both the headend and branch routers. The branch router can either have a static public interface IP address or one that is obtained dynamically from the service provider.

The routing control plane uses a dynamic IGP routing protocol such as EIGRP or OSPF over the VPN tunnels between headend and branch routers.

Topology

In a p2p GRE over IPsec design, only the following topologies are possible:

- Hub-and-spoke
- Partial mesh
- Full mesh

For all topologies listed above, administrative configuration is required. Unfortunately, there are no automatic configuration methods available for configuring the p2p GRE tunnel interfaces in Cisco IOS.

Hub-and-spoke topologies are the most common topologies in a p2p GRE over IPsec design. These topologies are the most scalable and predominately mimic traditional Layer 2 leased line, Frame Relay, or ATM hub-and-spoke networks.

Although partial mesh topologies are available, they are limited by both the routing protocol and the possibility of a dynamic public IP address. Configuring a partial mesh topology within a p2p GRE over IPsec design requires obtaining static public IP addresses for the branch routers that peer between each another.

Full mesh topologies are available as well and have the same limitations as partial mesh topologies. However, considering the administrative overhead involved, a full mesh topology is not recommended in a p2p GRE over IPsec design. If a full mesh topology is required, you should consider a DMVPN spoke-to-spoke topology, as outlined in the *Dynamic Multipoint VPN (DMVPN) Design Guide*, which is available at the following URL: <http://www.cisco.com/go/srnd>.

Headend System Architectures


The following two headend system architectures are described in this design guide:

- Single Tier Headend Architecture—Incorporates both the p2p GRE and crypto functions onto a single routing processor.
- Dual Tier Headend Architecture—Splits the p2p GRE and crypto functions onto two different routing processors.

Single Tier Headend Architecture

Figure 2-1 shows a Single Tier Headend Architecture for the p2p GRE over IPsec design.

Figure 2-1 p2p GRE over IPsec—Single Tier Headend Architecture



	Headend	Branch
Routing Control Plane	Dynamic Routing	Dynamic Routing
GRE Control Plane	Point-to-Point GRE	Point-to-Point GRE
IPSec Control Plane	Dynamic or Static Crypto Map	Static Crypto Map
	DPD	DPD
		Tunnel Protection

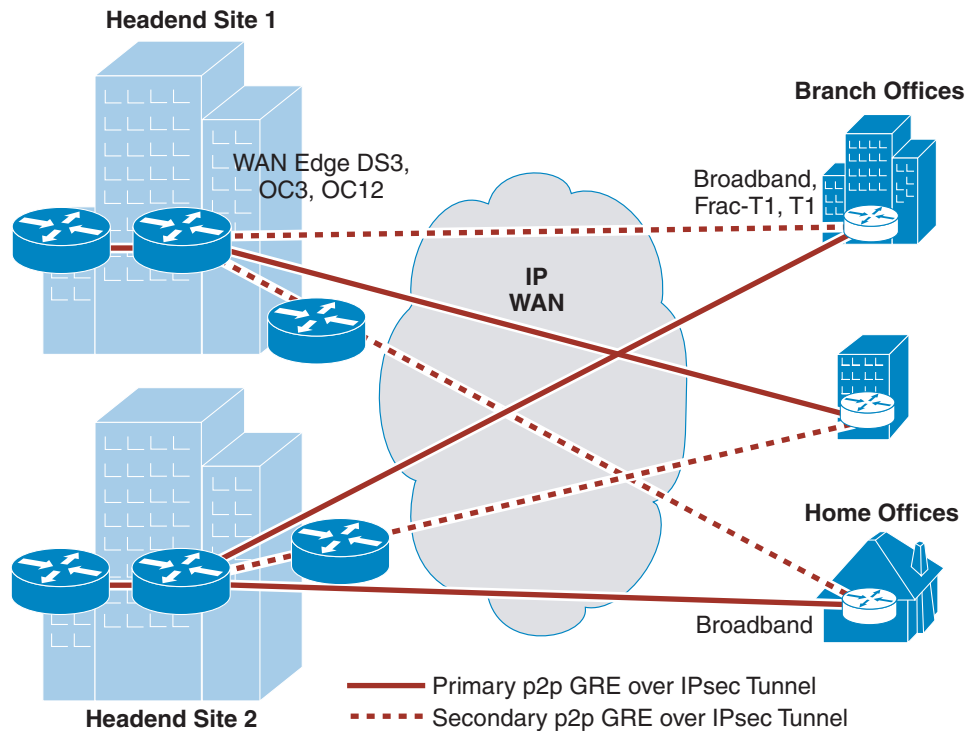
148878

The Single Tier Headend Architecture incorporates all three of the control planes shown in Figure 2-1 into a single routing processor. This architecture impacts scalability, where the central CPU becomes the gating factor.

Dual Tier Headend Architecture

Figure 2-2 shows a Dual Tier Headend Architecture for the p2p GRE over IPsec design.

Figure 2-2 p2p GRE over IPsec—Dual Tier Headend Architecture



	Headend		Branch	
Routing Control Plane	Dynamic Routing		Dynamic Routing	
GRE Control Plane	Point-to-Point GRE		Point-to-Point GRE	
IPsec Control Plane	Dynamic Crypto Map	DPD	Static Crypto Map	DPD

148879

The Dual Tier Headend Architecture incorporates the three control planes shown in Figure 2-2 into two routing processors. Both the routing and GRE control planes are housed on one routing process, while the IPsec control plane is housed on another. The reason for separating the functionality is to provide the best scalable solution given various platform limitations; specifically, CPU dependencies and resiliency.

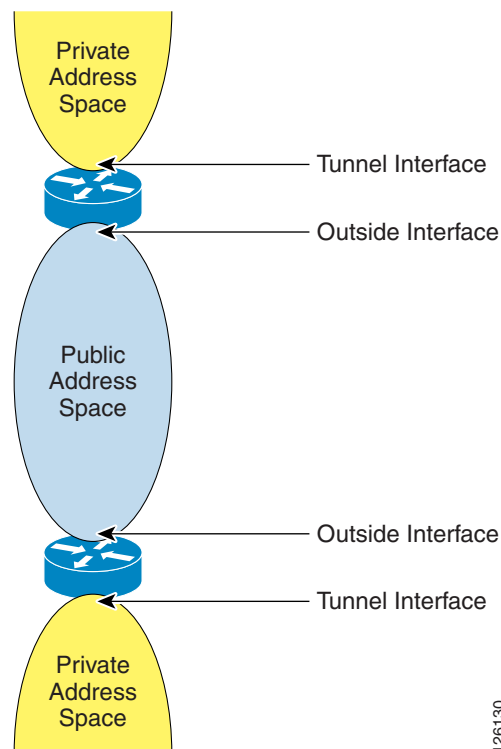
IP Addressing

Proper address summarization is highly recommended because it accomplishes the following:

- Conserves router resources, making routing table sizes smaller
- Saves memory in routers
- Eases troubleshooting tasks
- Simplifies the configuration of routers in IPsec networks

Although it is generally understood that VPNs are used for secure communications across a shared infrastructure (such as the Internet), make sure to distinguish between the enterprise addressing space, sometimes referred to as the private or inside addresses; and the infrastructure addressing space, also referred to as the service provider, public, or outside addresses. (See [Figure 2-3](#).)

Figure 2-3 Private and Public Address Spaces



126130


In most p2p GRE over IPsec VPN designs, the outside interface of the router is addressed in the infrastructure (or public) address space assigned by the service provider, while the tunnel interface belongs to the enterprise private network address space. In a static p2p GRE over a static IPsec configuration, the tunnel interfaces are sourced and destined to the public addresses. However, in the dynamic crypto peer address and static p2p GRE configuration, the branch router crypto IP address is dynamically obtained. For configuration details, see [Static p2p GRE over IPsec with a Branch Dynamic Public IP Address Case Study](#), page 5-1.

Generic Route Encapsulation

Although IPsec provides a secure method for tunneling data across an IP network, it has limitations. IPsec does not support IP broadcast or IP multicast, preventing the use of protocols that rely on these features, such as routing protocols. IPsec also does not support the use of multiprotocol traffic.

Generic Route Encapsulation (GRE) is a protocol that can be used to “carry” other passenger protocols, such as IP broadcast or IP multicast, as well as non-IP protocols. (See [Figure 2-4](#).)

Figure 2-4 GRE as a Carrier Protocol of IP



Using GRE tunnels in conjunction with IPsec provides the ability to run a routing protocol, IP multicast (IPmc), or multiprotocol traffic across the network between the headend(s) and branch offices.

GRE also enables private addressing. Without a tunnel protocol running, all end stations are required to be addressed with registered IP addresses. By encapsulating the IP packet in a tunneling protocol, private address space can be used.

With the p2p GRE over IPsec solution, all traffic between sites is encapsulated in a p2p GRE packet before the encryption process, simplifying the access control list used in the crypto map statements. The crypto map statements need only one line permitting GRE (IP Protocol 47).

GRE Keepalives

Beginning in Cisco IOS 12.2(8)T, the GRE keepalive feature is available for use on tunnel interfaces. This functionality allows the line protocol of the tunnel interface to track the reachability between the two tunnel endpoints. Beginning in Cisco IOS 12.2(11)T, the GRE keepalives are marked as DSCP value CS6.

If GRE keepalives are sent and acknowledged by the remote router, the line protocol is UP. If successive GRE keepalives are not acknowledged, based on the configured interval and number of retries, the tunnel line protocol is marked DOWN.

If the network manager has configured a routing protocol for the tunnel, the routing protocol (RP) hello packets provide at Layer 3 a similar function to the GRE keepalive. However, it may be desirable from a network management standpoint to be able to generate a Simple Network Management Protocol (SNMP) trap when the p2p GRE interface line protocol goes down. This is an example where running both Layer 2 (GRE) and Layer 3 (RP hello) is advantageous.

There are advantages to eliminating the routing protocol and relying on the GRE keepalive to verify connectivity. If the branch router is a stub network with no need for full routing information, a default route can be configured to the tunnel interface on the branch router, and the headend router can redistribute a static route using the tunnel interface name as the next hop. If the GRE keepalives are lost, the line protocol goes DOWN, and the redistributed route is withdrawn from the routing table and advertisements to other RP neighbors.

This reduces the number of RP peers the headend router must maintain, and the branch router configuration is simplified because no RP must be configured. Network stability and performance may be enhanced by reducing the CPU required for the overhead function of maintaining RP neighbors, and instead using those CPU cycles for packet switching.

Using a Routing Protocol across the VPN

This design recommends the use of a routing protocol to propagate routes from the headend to the branch offices. Using a routing protocol has several advantages over the current mechanisms in IPsec Direct Encapsulation alone.

In a VPN, routing protocols provide the same level of benefits as compared to a traditional network, including the following:

- Network topology information
- Topology change notification (such as when a link fails)
- Remote peer status

Several routing protocols are candidates for operation over a p2p GRE over IPsec VPN, including EIGRP and OSPF. Designs presented in this design guide use EIGRP as the routing protocol because EIGRP was used during the scalability tests conducted. EIGRP is recommended as the routing protocol because of its conservative use of router CPU and network bandwidth as well as its quick convergence times. EIGRP also provides a range of options for address summarization and default route propagation.

Other routing protocols, such as OSPF, have been verified in designs, but are not discussed in this design guide.

Routing protocols do increase the CPU utilization on a network device, and this impact must be considered when sizing those devices.

Route Propagation Strategy

There are a number of approaches to propagating routes from the headend to the branch offices. For this design, the recommended approach is for each headend router to advertise either a default route or summary routes down each of the tunnels, with a preferred routing metric for the primary path. With this in mind, each of the branch office routers need to add a static host route for each of the headend peer (primary and secondary) IP addresses, with a next hop destined for their respective ISP IP address. The purpose for the static host routes is to avoid recursive routing through the p2p GRE tunnel. Recursive routing occurs when a route to the p2p GRE tunnel source outside IP address of the opposing router is learned via a route with a next hop of the inside IP address of the opposing p2p GRE tunnel. This breaks the tunnel because it causes the p2p GRE encapsulated packet to be routed into its own p2p GRE tunnel instead of being routed directly.

Crypto Considerations

The use of crypto is imperative to the p2p GRE over IPsec design because it provides the secure channel between the headend and branch routers. The p2p GRE tunnel is encrypted inside the crypto tunnel. For specific crypto considerations, see the *IPsec Direct Encapsulation Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

IPsec Tunnel versus Transport Mode

Integrating p2p GRE with either IPsec tunnel mode or transport mode has been debated. Tunnel mode adds an additional 20 bytes to the total packet size. Either tunnel or transport mode work in a p2p GRE over IPsec implementation; however, several restrictions with transport mode should be considered. If the crypto tunnel transits either a Network Address Translation (NAT) or Port Address Translation (PAT) device, tunnel mode is required. In addition, this design guide shows configuration examples for implementing p2p GRE over IPsec where the p2p GRE tunnel endpoints are different than the crypto tunnel endpoints. Tunnel mode is also required in these cases.

Dead Peer Detection

Dead Peer Detection (DPD) is a relatively new Cisco IOS feature that is actually an enhancement of the ISAKMP keepalives feature. DPD operates by sending a hello message to a crypto peer from which it has not received traffic during a specified configurable period. If normal IPsec traffic is received from a crypto peer and decrypted correctly, that crypto peer is assumed alive, no hello message is sent, and the DPD counter for that crypto peer is reset. This results in lower CPU utilization than that which would have occurred with ISAKMP keepalives.

In the event that no traffic is received during the specified period, an ISAKMP R_U_THERE message is sent to the other crypto peer. If no response is received after the specified number of tries, the connection is assumed dead, and the IPsec tunnel is disconnected. This feature is vital to prevent black-holing traffic, in the event that the Security Association (SA) database of one side is cleared manually or by reboot. DPD is both a headend and branch technology and should be configured on both sides of a VPN tunnel.

DPD should always be configured, even when GRE keepalives or a routing protocol are used.

Configuration and Implementation

The configuration issues defined in this chapter are specific to VPN implementation for the p2p GRE over IPsec design topology. It is presumed that the reader is reasonably familiar with standard Cisco configuration practices at the command-line interface (CLI) level.

All configuration examples shown are for IPsec in tunnel mode. Also, all references to private or public IP addresses correlate to [IP Addressing, page 2-5](#).

For more details and a step-by-step instruction, see the following URL:

http://www.cisco.com/en/US/partner/tech/tk583/tk372/tsd_technology_support_protocol_home.html

ISAKMP Policy Configuration

There must be at least one matching ISAKMP policy between two potential crypto peers. The sample configuration below shows a policy using Pre-Shared Keys (PSK) with 3DES as the encryption algorithm. There is a default ISAKMP policy that contains the default values for the encryption algorithm, hash method or Hashed Method Authentication Code (HMAC), Diffie-Hellman group, authentication type, and ISAKMP SA lifetime parameters. This is the lowest priority ISAKMP policy.

When using PSK, Cisco recommends that wildcard keys not be used. However, when implementing a p2p GRE over IPsec design using an IP address obtained dynamically, the use of a wildcard PSK or Public Key Infrastructure (PKI) on the headend router is required. Instead, the example shows two keys configured for two separate crypto peers. The keys should be carefully chosen; “bigsecret” is used only as an example. The use of alphanumeric and punctuation characters as keys is recommended.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.161.2
```

Branch router:

```
interface Serial0/0
ip address 192.168.161.2 255.255.255.0
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.251.1
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, a wildcard PSK or PKI must be used on the crypto headend router.

For more information regarding configuring ISAKMP policies, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fipsencr/srfike.htm#wp1017989

Dead Peer Detection Configuration

An enhancement to the **crypto isakmp keepalive** command has changed the way that ISAKMP keepalives work, creating the feature known as Dead Peer Detection (DPD). DPD no longer automatically sends hello messages to the ISAKMP peer if live traffic has been received from that peer within a specified period. The first variable in the **crypto isakmp keepalive** command is the number of seconds that the peer waits for valid traffic from its crypto neighbor. If no traffic has been received, the second variable is the number of seconds between retries. This scheme helps conserve router CPU by not sending the keepalive messages if a router has just received valid traffic.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
```

```

ip address 192.168.251.1 255.255.255.0
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.161.2
crypto isakmp keepalive 10
!

```

Branch router:

```

interface Serial0/0
ip address 192.168.161.2 255.255.255.0
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp keepalive 10
!

```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, a wildcard PSK or PKI must be used on the crypto headend router.

IPsec Transform and Protocol Configuration

The transform set must match between the two IPsec peers. The transform set names are locally significant only. However, the encryption algorithm, hash method, and the particular protocols used (ESP or AH) must match. You may also configure data compression here but it is not recommended on peers with high speed links. There can be multiple transform sets for use between different peers, with the strongest match being negotiated.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```

interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.161.2
crypto isakmp keepalive 10
!
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac

```

Branch router:

```

interface Serial0/0
ip address 192.168.161.2 255.255.255.0
!

```

```
crypto isakmp policy 10
  encr 3des
  authentication pre-share
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp keepalive 10
!
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, a wildcard PSK or PKI must be used on the crypto headend router.

For more information on transform sets and configuring crypto maps, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fipsencr/srfipsec.htm#xtocid105784

Access Control List Configuration for Encryption

The access control list entries defining the traffic to be encrypted should be mirror images of each other on the crypto peers. If access control list entries include ranges of ports, a mirror image of those same ranges must be included on the access control lists of the remote peer. The addresses specified in these access control lists are independent of the addresses used by the crypto peers. This example specifies the IP protocol GRE on both the source and destination parts of the access control list. All traffic encapsulated in the p2p GRE packets is protected.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```
ip access-list extended vpn-static1 permit gre host 192.168.251.1 host 192.168.1.2
```

Branch router:

```
ip access-list extended vpn-static2 permit gre host 192.168.1.2 host 192.168.251.1
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router. However, note that the p2p GRE headend source and destination public IP addresses are different from the crypto headend. The crypto ACL needs to match the p2p GRE tunnel endpoints.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, the headend crypto ACL is *not* required. The headend router uses a dynamic crypto map that dynamically creates its crypto ACL from the incoming branch router crypto ACL. The branch router ACL is identical to the configuration example above.

Crypto Map Configuration

The crypto map entry ties together the crypto peers, the transform set used, and the access control list used to define the traffic to be encrypted. The crypto map entries are evaluated sequentially.

In the example below, the crypto map name “static-map” and crypto map numbers (for example, “10” and “20”) are locally significant only. The first statement sets the IP address used by this peer to identify itself to other crypto peers in this crypto map. This address must match the set peer statement in the crypto map entries of the remote crypto peers. This address also needs to match the address used with any PSK the remote peers might have configured. The IPsec mode defaults to tunnel mode.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
!
crypto map static-map local-address FastEthernet1/0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.161.2
  set transform-set vpn-test
  match address vpn-static1
```

Branch router:

```
interface Serial0/0
ip address 192.168.161.2 255.255.255.0
!
crypto map static-map local-address Serial0/0
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address vpn-static2
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router

The following configuration example shows a dynamic public IP address on the branch router with a static public IP address on the headend router for the crypto peers for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
!
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
!
!
crypto map dynamic-map local-address FastEthernet1/0
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
```

Branch router:

```
interface Serial0/0
```



```
ip address dhcp
!
crypto isakmp key bigsecret address 192.168.251.1
!
crypto map static-map local-address Serial0/0
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address vpn-static2
```

On the headend router, a dynamic crypto map is used with a wildcard PSK to allow a crypto peer with the public dynamically served IP address of the branch router.

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.

For a more complete description of the various crypto configuration commands, see the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgr/fsecur_r/fipseocr/srfipseocr.htm

Applying Crypto Maps

In releases before Cisco IOS Release 12.2(13)T, the crypto maps must be applied to both the physical interface and the logical interfaces, such as the p2p GRE tunnel interfaces. As of Cisco IOS Release 12.2(13)T (assumed in the example below), the crypto map is applied only to the physical interface, not to the logical interface.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the crypto peer for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
crypto map static-map
```

Branch router

```
interface Serial0/0
ip address 192.168.161.2 255.255.255.0
crypto map static-map
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.

The following configuration example shows a public dynamic IP address on the branch router with a static public IP address on the headend router for the crypto peers for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0
crypto map dynamic-map
```

Branch router:

```
interface Serial0/0
 ip address dhcp
 crypto map static-map
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the crypto headend router.

Tunnel Interface Configuration—Branch Static Public IP Address

This section shows the tunnel interface configurations using a branch static public IP address.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the p2p GRE tunnel for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface Tunnel0
 bandwidth 1536
 ip address 10.62.1.193 255.255.255.252
 tunnel source 192.168.251.1
 tunnel destination 192.168.161.2
```

Branch router:

```
interface Tunnel0
 bandwidth 1536
 ip address 10.62.1.194 255.255.255.252
 tunnel source 192.168.161.2
 tunnel destination 192.168.251.1
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the p2p GRE headend router. The p2p GRE headend router has a different static public IP address than the crypto headend router.

Tunnel Interface Configuration—Branch Dynamic Public IP Address

This section shows the tunnel interface configurations using a branch dynamic public IP address.

The following configuration example shows a dynamic public IP address on the branch router with a static public IP address on the headend router for the p2p GRE tunnel for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
 ip address 192.168.251.1 255.255.255.0
 !
interface Tunnel0
 bandwidth 1536
 ip address 10.62.1.193 255.255.255.252
```

```

tunnel source 192.168.251.1
tunnel destination 10.62.1.255
!
ip route 10.62.1.255 255.255.255 192.168.251.2

```

Branch router:

```

interface Serial10/0
 ip address dhcp
!
interface Loopback0
 ip address 10.62.1.255 255.255.255.255
!
interface Tunnel0
 bandwidth 1536
 ip address 10.62.1.194 255.255.255.252
 tunnel source 10.62.1.255
 tunnel destination 192.168.251.1

```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the p2p GRE headend router. The p2p GRE headend router has a different static public IP address than the crypto headend router. The static host route of the p2p GRE headend router to the Loopback0 IP address of the branch router may not be required because the p2p GRE headend router sends all traffic to the crypto headend router.

For more detailed information, see [Static p2p GRE over IPsec with a Branch Dynamic Public IP Address Case Study, page 5-1](#).

GRE Keepalive Configuration

This section shows a sample headend and branch configuration using GRE keepalives.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the p2p GRE tunnel for either a Single or Dual Tier Headend Architecture:

Headend router:

```

interface Tunnel0
 ip address 10.62.1.193 255.255.255.252
 keepalive 10 3
!
ip route 10.62.1.0 255.255.255.0 10.62.1.194

```

Branch router:

```

interface Tunnel0
 ip address 10.62.1.194 255.255.255.252
 keepalive 10 3
!
ip route 10.0.0.0 255.0.0.0 10.62.1.193

```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.

- In a Dual Tier Headend Architecture, the configuration above is applied to the p2p GRE headend router. The p2p GRE headend router has a different static public IP address than the crypto headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, the configuration above is the same.

GRE keepalives are a trigger mechanism to cause the line protocol to be changed from an UP/UP to an UP/DOWN state during a failure event. A floating static route can be used in place of a routing protocol on the branch router. In the headend router, a routing protocol may be required to redistribute the static routes into the campus network topology.

Routing Protocol Configuration

This section shows a sample headend and branch configuration using EIGRP as the routing protocol.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the p2p GRE tunnel for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet0/0
 ip address 10.57.1.1 255.255.255.0
!
interface Tunnel0
 ip address 10.62.1.193 255.255.255.252
!
router eigrp 10
 network 10.0.0.0
 no auto-summary
```

Branch router:

```
interface FastEthernet0/0
 ip address 10.62.1.1 255.255.255.128
!
interface Tunnel0
 ip address 10.62.1.194 255.255.255.252
!
router eigrp 10
 network 10.0.0.0
 no auto-summary
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the p2p GRE headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, the configuration above is the same.

Route Propagation Configuration

This section shows a sample headend and branch configuration using EIGRP as the routing protocol redistributing a static route into the EIGRP routing process.

The following configuration example shows a static public IP address on the branch router with a static public IP address on the headend router for the p2p GRE tunnel for either a Single or Dual Tier Headend Architecture:

Headend router:

```
interface FastEthernet1/0
 ip address 192.168.251.1 255.255.255.0
!
router eigrp 10
 network 10.0.0.0
 no auto-summary
 redistribute static metric metric 10000 10 255 1 1500
!
ip route 0.0.0.0 0.0.0.0 192.168.251.2
```

Branch router:

```
interface Serial10/0
 ip address dhcp
!
router eigrp 10
 network 10.0.0.0
 no auto-summary
!
ip route 192.168.251.1 255.255.255.255 dhcp
```

Note the following:

- In a Single Tier Headend Architecture, the configuration above is applied to the headend router.
- In a Dual Tier Headend Architecture, the configuration above is applied to the p2p GRE headend router.
- In either headend architecture implementing a static p2p GRE over IPsec with a branch dynamic public IP address, the configuration above is the same.

In the above example, a default route is being redistributed into EIGRP AS 10 on the headend router and then advertised to the branch router with an administrative distance (AD) of 90. Considering that the branch router has a default route learned via DHCP with an AD of 254, recursive routing *must* be taken into account. To avoid recursive routing on the branch router, a static host route for the crypto peer address is added to the configuration to ensure that the outside of the tunnel is routed directly to the ISP instead of inside the p2p GRE tunnel.

High Availability

High Availability (HA) provides network resilience and availability in the event of a failure. This section provides some designs for highly available p2p GRE over IPsec VPNs. HA is covered in much more depth in the *IPsec VPN Redundancy and Load Sharing Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.


Common Elements in all HA Headend Designs

To provide a level of resiliency in the VPN design, Cisco recommends that at least two tunnels be configured on each branch. Each branch router should have a tunnel to a primary headend, and an alternate tunnel to a secondary headend. Under normal operating conditions, both the primary and secondary tunnels have routing protocol neighbors established. The routing protocol maintains both paths, with the secondary tunnel being configured as a less preferred path.

A common concern in all HA headend resilient designs is the number of RP neighbors. Many redundant neighbor relationships increase the time required for routing convergence.

Figure 2-5 shows a typical HA scenario.

Figure 2-5 Branch Router Connected via p2p GRE over IPsec to More Than One Headend Device



If a failure occurs at one of the headend devices, the routing protocol detects that the route through the primary tunnel is no longer valid and, after convergence, the route through the secondary tunnel is used. When the primary is available again, traffic is routed back to the primary tunnel because it is the preferred route in the routing metrics.

The headend resiliency design presented here allows for failure of a single headend device, with proper failover to surviving headends. The typical branch router has two or more tunnel interfaces to two or more VPN headends; the site location of these is an architectural decision of the HA strategy.

In all HA architectures, all tunnels from the branch to the headend routers are up. The routing protocol determines which tunnel is passing user traffic. The different paths in this design are configured with slightly different metrics to provide preference between the tunnels. The routing metric should be consistent both upstream and downstream to prevent asymmetric routing.

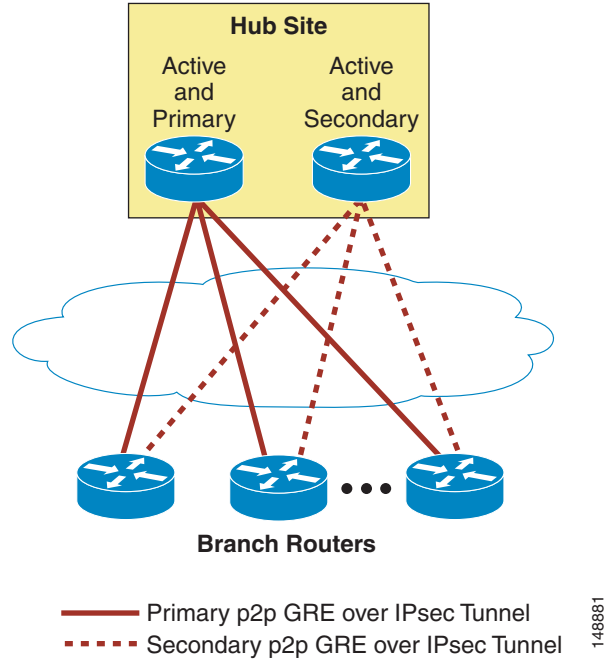
The following sections describe some commonly used architectures in the headend HA design.

1+1 (Active-Standby) Failover Headend Resiliency Design

In a 1+1 failover, each primary headend is paired with a standby headend. The primary headend is passing user traffic, while the standby headend is maintaining p2p GRE tunnels and routing neighbors. The routing protocol determines which p2p GRE tunnel is the active path for user traffic.

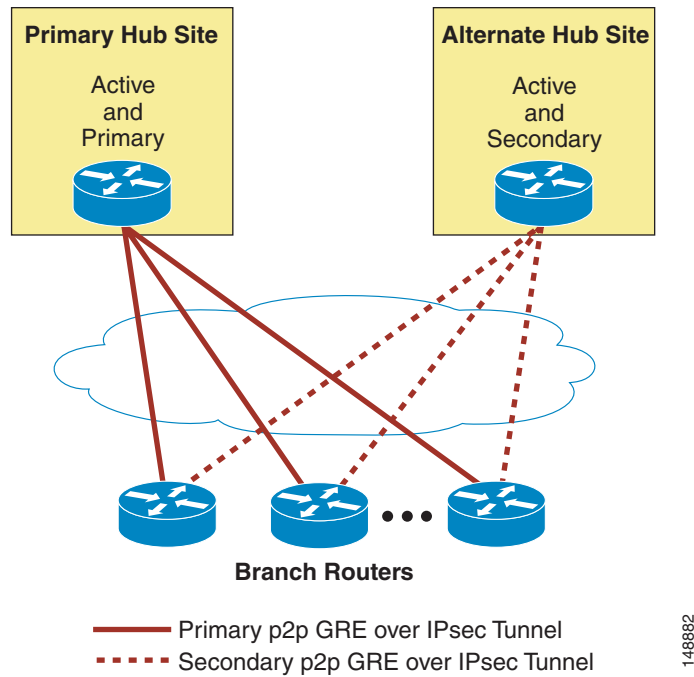
1+1 failover headends may be deployed in one site or in different sites. Figure 2-6 and Figure 2-7 show these topologies.

Figure 2-6 *Box Redundancy—HA p2p GRE over IPsec with Two Crypto Headends in One Hub Site*



It may also be necessary in the customer strategy to have headend devices geographically dispersed. One such design is shown in Figure 2-7:

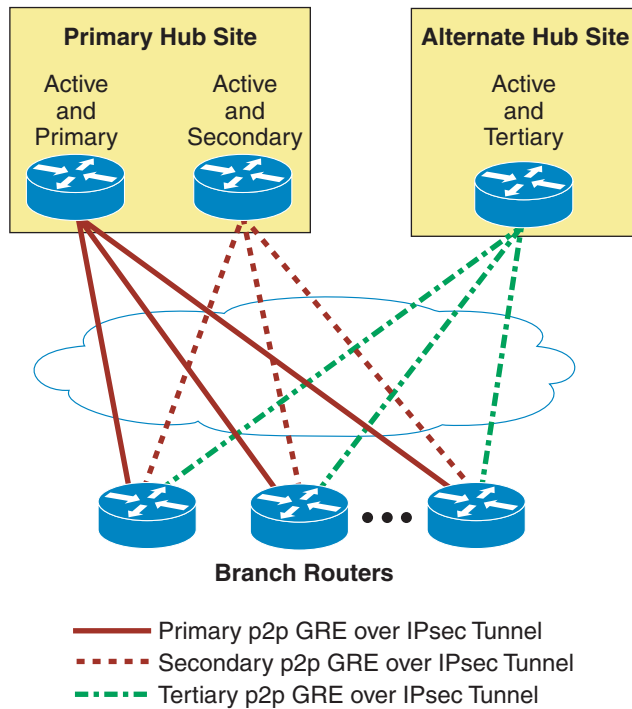
Figure 2-7 *Site Redundancy—HA p2p GRE over IPsec with One Crypto Headend in Each Hub Site*



In this design example, each remote router has a primary p2p GRE over IPsec tunnel to a headend at the primary site, as well as a secondary tunnel to a different headend at a different site (site redundancy).

A network manager can also do a combination of both box and site redundancy on a respective branch at the same time. [Figure 2-8](#) shows this topology.


Figure 2-8 Combined Redundancy—HA p2p GRE over IPsec with Multiple Crypto Headends in Various Locations



Load Sharing with Failover Headend Resiliency Design

Another possibility for a headend redundancy design is shown in [Figure 2-9](#).

Figure 2-9 Load Sharing with Failover HA



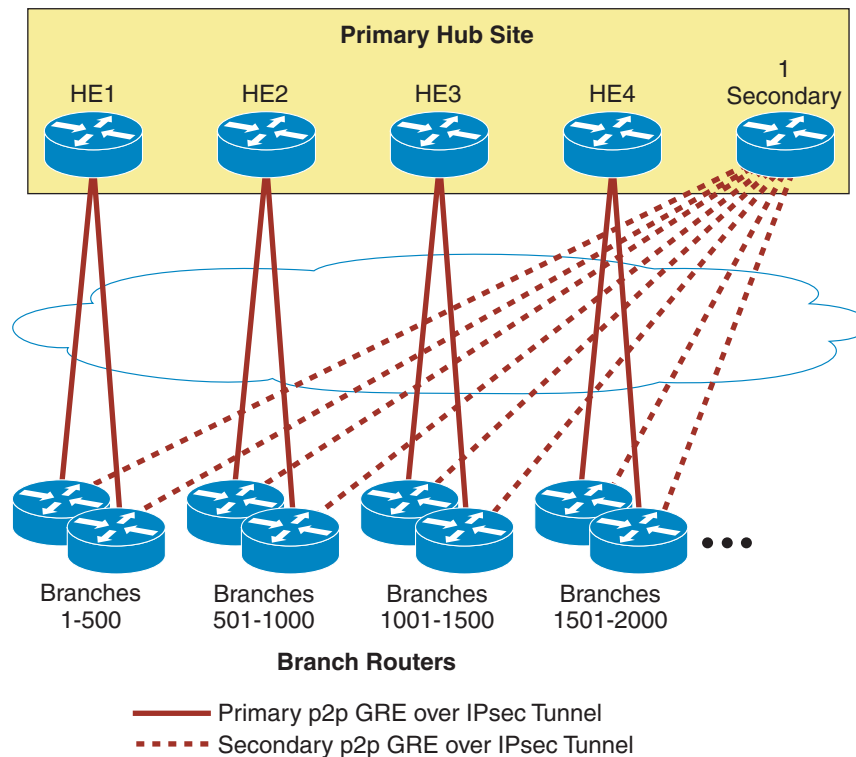
In this design, each branch has a primary path, which is used to pass traffic under normal conditions. Each branch has a secondary path in the event of a failover occurrence with the primary headend. This failover strategy uses a manually configured distribution across the headend devices. In [Figure 2-9](#), each headend carries approximately one-third of the user traffic, as well as being a secondary headend for another one-third of the user traffic in the event of a failure. The network manager must take care to properly scale the amount of tunnels and traffic to a particular headend system to ensure that any headend device can carry its normal load, as well as its failover load, and remain at a reasonable CPU and pps level for the given platform.

A network manager may add headend devices to this series. This addition requires manually changing the distribution, and requires modification to both the branch router configurations as well as the affected headends.

N+1 Failover Architecture

In an N+1 failover, each group of branches has a primary path to their respective headend system and a secondary path to the one and only one common secondary system. [Figure 2-10](#) shows this topology.

Figure 2-10 N+1 Failover HA



148884

This failover architecture is not recommended because the secondary (standby) system is required to maintain p2p GRE over IPsec tunnels and routing neighbors to all the branches for which it is a secondary.

Using [Figure 2-10](#) as an example, scalability concerns illustrate why the topology can exceed the following limitations:

- The number of recommending routing neighbors on the secondary (should not exceed the RP recommendations)
- The limitation of the CLI in Cisco IOS on the number of tunnel interfaces that can be configured and supported in one system (platform-dependant)
- The limit of the number of IPsec peers that one system can effectively maintain and re-key
- The pps rate of a failed primary to the secondary (with the addition of the previous three issues above) may oversubscribe the single secondary

Dual Tier Headend Architecture Effect on Failover

The architectures shown in the previous sections have been Single Tier Headend Architectures (crypto, GRE, and RP all on one headend system). If a Dual Tier Headend Architecture is implemented, the crypto functionality is separated from the GRE and RP functions. The crypto failover portion now has more failover options (see Section 4.3 of the *IPsec Direct Encapsulation Design Guide* at the following URL: <http://www.cisco.com/go/srnd>).

The following p2p GRE and RP strategies are still valid architectures for the traffic failover:

- 1+1 failover (box, site, or combined)
- Load sharing with failover

QoS

To support latency-sensitive traffic applications, it may be necessary to configure QoS. QoS and IPsec have been integrated as part of the Cisco Voice and Video Enabled IPsec VPN (V3PN) technology.

For more information, see the following documents:

- *Voice and Video VPN (V3PN): QoS and IPsec Design Guide*— <http://www.cisco.com/go/srnd>
- *Enterprise QoS Solution Reference Network Design Guide*— http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf.

IP Multicast

Scalability testing with IP multicast and IPsec encryption indicates that there are issues with packet loss, because of the instant replication of many packets. IP multicast replication happens at a single moment in time. The replication occurs before encryption, meaning that the crypto cards or engines in the various platforms can be overwhelmed if a large number of spokes are joined to the same IP multicast stream.

For example, consider a design using the Cisco Catalyst 6500 with VPN SPA, and configuring 1000 p2p GRE over IPsec tunnels to branch offices. If each branch office is joined to a single IP multicast stream, the VPN SPA must replicate each IP multicast packet 1000 times, one per VPN tunnel. Assuming the Sup720 can sustain the replication speed of the stream, many packets (up to 1000) arrive at the input queue of the VPN SPA, causing overruns or dropped packets.

For appropriate scalable designs if the customer has multicast requirements, see the *Multicast over IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Interactions with Other Networking Functions

This section describes other networking functions such as PAT, DHCP, and firewall considerations that apply to designing a p2p GRE over IPsec design.

Network Address Translation and Port Address Translation

Although NAT and PAT can result in an added layer of security and address conservation, they both present challenges to the implementation of an IPsec VPN. ISAKMP relies on an individual IP address per crypto peer for proper operation. PAT works by masquerading multiple crypto peers behind a single IP address.

The IPsec NAT Traversal feature (NAT-T) introduces support for IPsec traffic to travel through NAT or PAT devices by encapsulating both the IPsec SA and the ISAKMP traffic in a UDP wrapper. NAT-T was first introduced in Cisco IOS version 12.2(13)T, and is auto-detected by VPN devices. There are no configurations steps for a Cisco IOS router running this release or later because it is enabled by default as a global command. The NAT-T feature detects a PAT device between the crypto peers and negotiates NAT-T if it is present.

For more details on IPsec NAT-T, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ftipsnat.htm>

Dynamic Host Configuration Protocol

For a host at a remote site to be able to use a DHCP server over an IPsec tunnel at a central site, an IP helper address must be configured on the router interface associated with the host.

One drawback of this approach is that if connectivity to the central site is lost, a host at a remote site may not receive or renew an IP address. The inability to receive an IP address results in the host being unable to communicate to the local network.

A Cisco IOS router can be configured as a DHCP server. Using the router as a stand-alone DHCP server is recommended for branch offices with no redundant links.

Firewall Considerations

This section describes the various firewall considerations when implementing a p2p over GRE design.

Headend or Branch

Depending on the crypto and p2p GRE headend or branch placements, the following protocols and ports are required to be allowed:

- UDP Port 500—ISAKMP as source and destination
- UDP Port 4500—NAT-T as a destination
- IP Protocol 50—ESP
- IP Protocol 51—AH (if AH is implemented)
- IP Protocol 47—GRE (if GRE traverses the firewall post decryption)
- Any potential end user traffic—If GRE does not traverse the firewall post encapsulation

Network location of the crypto headend in relation to the headend firewall(s) impacts both the accessibility and performance of the both systems. The network manager must ensure that all firewalls are properly configured to allow the tunnel traffic bi-directionally. The crypto headend must be accessible to the branch router.

Firewall Feature Set and Inbound ACL

Before Cisco IOS version 12.3(8)T, packets received on an interface with an inbound ACL and a crypto map were checked by the inbound ACL twice, before decryption, and as clear-text following decryption. The Crypto Access Check on Clear-Text Packets feature removes the checking of clear-text packets that go through the IPsec tunnel just before or just after decryption.

Double ACL Check Behavior (Before 12.3(8)T)

If the enterprise security policy does not permit split tunnel, and the branch requires Internet access through the IPsec tunnel, the remote routers must also be configured to permit specified TCP and UDP traffic through the inbound access control list when the connection is initiated from within the remote router subnet.

To allow Internet access in non-split tunnel configurations, use Context-Based Access Control (CBAC) in conjunction with the inbound access control list:

```
ip inspect name CBAC tcp
ip inspect name CBAC udp
ip inspect name CBAC ftp
ip inspect name CBAC sip
!
interface Ethernet 0
description Inside
 ip address 10.81.7.1 255.255.255.248
!
interface Ethernet 1
description Outside
 ip address dhcp
 ip access-group INPUT_ACL in
 ip inspect CBAC out
!
ip access-list extended INPUT_ACL
 permit udp x.x.x.16 0.0.0.15 any eq isakmp
 permit udp x.x.x.16 0.0.0.15 any eq non500-isakmp
 permit esp x.x.x.16 0.0.0.15 any
 remark ! Enterprise Address space
 permit ip 10.0.0.0 0.255.255.255 10.81.7.0 0.0.0.7
 permit udp any any eq bootpc
 permit udp x.x.x.40 0.0.0.1 eq ntp any
 permit tcp x.x.0.0 0.0.15.255 any eq 22
 permit icmp any any
 deny ip any any
end
```

Crypto Access Check on Clear-Text Packets Feature (12.3(8)T and Later)

The Crypto Access Check on Clear-Text Packets feature removes the checking of inbound, just-decrypted clear-text packets against the outside interface inbound ACL.

When upgrading Cisco IOS to a version that supports this feature, the following statement should be removed from the **ip access-list extended INPUT_AC** command, and the **ip inspect CBAC in** command can be removed from interface Ethernet 0:

```
! Enterprise Address space
 permit ip 10.0.0.0 0.255.255.255 10.81.7.0 0.0.0.7
```

If checking the decrypted clear-text packets against an ACL is desired, that function is now configured inside the crypto map global configuration.

For more information on Crypto Access Check on Clear-Text Packets, see the following URL:

http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5207/products_feature_guide09186a008022c2a5.html

Common Configuration Mistakes

The following sections outline some common mistakes and problems encountered when configuring p2p GRE over IPsec.

Crypto Peer Address Matching using PSK

The IP address used as the crypto source address must match the address configured as the destination address on the crypto peer, and vice-versa. Unless the address is configured specifically, the address of the outgoing interface is used as the crypto peer address, thus causing the crypto peer to die at ISAKMP negotiation.

Transform Set Matches

At least one matching IPsec transform set must be configured between two crypto peers. When specifying a particular strength of encryption algorithm, a similar strength encryption algorithm should also be configured. Failure to do so can weaken the encryption strength of the entire solution.

ISAKMP Policy Matching

There is a default ISAKMP policy present in all Cisco IOS devices. This default is encryption DES, HMAC of SHA, IKE authentication of RSA signature, and DH group 1. If a stronger ISAKMP policy is desired, both sides must support that policy.

It is common, but not required, to use the same encryption level transform set and hash methods in ISAKMP policy and IPsec transform set.



Scalability Considerations

This chapter presents the steps to selecting Cisco products for a VPN solution, starting with sizing the headend, and then choosing Cisco products that can be deployed for headend devices. This chapter concludes with product sizing and selection information for branch devices.

General Scalability Considerations

This section provides general scalability considerations to assist with design requirements.

IPsec Encryption Throughput

The throughput capacity of the IPsec encryption engine in each platform (headend or branch) must be considered for scalable designs, because each packet that is encrypted must traverse through the encryption engine.

Encryption throughput must therefore consider bi-directional speeds. Several examples are shown in [Table 3-1](#) and [Table 3-2](#) for popular headend and branch connection speeds.

Table 3-1 Headend Connection Speeds

Connection Type	Speed (in Mbps)	Encryption throughput required (in Mbps)
T3/DS3	44.7	90.0
OC3	155.0	310.0
OC12	622.0	1250.0

Table 3-2 Branch Connection Speeds

Connection Type	Speed (in Mbps)	Encryption Throughput Required (in Mbps)
T1	1.5	3.0
2 x T1	3.0	6.0
T3/DS3	44.7	90.0
Broadband cable/DSL	384 Kbps uplink/ 2 Mbps downlink	2.4

In general, as throughput increases, the burden on router CPU also increases. However, with hardware-accelerated encryption available for all Cisco router products from the 871 through the 7600, impact on the main CPU is offloaded to the VPN hardware. Main router CPU processing still occurs, however, so higher throughput typically results in higher CPU utilization.

Packets Per Second—Most Important Factor

Although bandwidth throughput capacity must be considered, even more important is the packet rate for the connection speeds being terminated or aggregated.

In general, routers and encryption engines have upper boundaries for processing a given number of packets per second (pps). Size of packets used for testing and throughput evaluations can understate or overstate true performance. For example, if a router with a VPN module can handle 20 Kpps, then 100-byte packets lead to 16 Mbps throughput, while 1400-byte packets at the same packet rate lead to 224 Mbps.

Because of such a wide variance in throughput, pps is generally a better parameter than bits per second (bps) to determine router forwarding potential. Scalability of the headend is the aggregate forwarding potential of all branches that terminate a tunnel to that headend. Therefore, the aggregate pps from all branches impacts the pps rate of that headend.

Tunnel Quantity Affects Throughput

Although it is highly dependent on platform architecture, the overall throughput generally tends to decrease as tunnel quantities are increased. When a router receives a packet from a different peer than the one whose packet was just decrypted, a lookup based on the security parameters index of the new packet must be performed. The transform set information and negotiated key of the new packet is then loaded into the hardware decryption engine for processing. Traffic flowing on larger numbers of SAs tends to negatively affect throughput performance.

Increasingly, platforms with hardware-accelerated IPsec encryption are designed to offload tunnel processing overhead as well, resulting in more linear performance regardless of the number of tunnels. For example, the VPN SPA blade for the Cisco 7600 has relatively linear throughput regardless of whether the traffic load is offered on a few tunnels or several thousand. The Cisco VPN Services Adapter (VSA) for Cisco 7200 Series also incorporates these features and as such, performance is improved accordingly over a VAM2+ adapter.

GRE Encapsulation Affects Throughput

Router encryption throughput is affected by the configuration of GRE. In addition to the headers that are added to the beginning of each packet, these headers must also be encrypted. The GRE encapsulation process, when not hardware-accelerated, increases total CPU utilization. Total throughput in a p2p GRE over IPsec design results in a lower throughput than that of an IPsec Direct Encapsulation design.

Routing Protocols Affect CPU Overhead

CPU overhead is affected by running a routing protocol. Router processing of keepalives and maintenance of a routing table uses a finite amount of CPU time, which varies with the number of routing peers and the size of the routing table. The network manager should design the routing protocol based on well-known and accepted practices.

Headend Scalability

Headend devices are primarily responsible for the following:

- Terminating p2pGRE over IPsec tunnels from the branch routers
- Running a routing protocol inside the p2p GRE tunnels to advertise internal routes to the branches
- Providing redundancy to eliminate the possibility of a single point of failure

It is important to size the headend correctly before choosing the devices to deploy. This ensures that the overall network can support the intended (and possibly future) traffic profiles that the enterprise wants to run over the VPN.

The following critical factors must be considered when sizing the headend:

- How many branch offices need to be connected to the headend? This information provides the number of primary tunnels requiring aggregation.
- What is the expected traffic profile, including the average pps and bps throughput rates for each branch office? This information provides the aggregated data throughput required across the VPN.
- What is the headend connection speed?
- What is the high availability requirement for the design?
- What is the expected performance margin or target CPU utilization?

In general, it is recommended that headend devices be chosen so that CPU utilization does not exceed the target value. This recommendation is to ensure that the device has enough performance remaining to deal with various events that take place during the course of normal network operations, including network re-convergence in the event of a failure, re-keying IPsec SAs, and bursts of traffic.

However, keep in mind that the target value is just a guideline and your customer requirements should determine where to set the limits of operation. In addition, for some platforms such as the Catalyst 6500 and Cisco 7600 router, CPU is not an accurate reflection of performance limits. Other factors may limit the performance, such as backplane and packet switching speeds.

To provide an idea of design scalability and limits, several different design topologies and headend platforms have been evaluated under typical customer traffic profiles. These scalability test results are presented in [Chapter 4, “Scalability Test Results \(Unicast Only\).”](#)

The primary platform choices available today for headend routers are the following:

- Cisco 7200VXR with NPE-G1 and SA-VAM2+ encryption module
- Cisco 7200VXR with NPE-G2 and either a SA-VAM2+ or VPN Services Adapter (VSA)
- Cisco Catalyst 6500 or Cisco 7600 with Sup720, SSC400, and VPN SPA

The Cisco 7301 router is also an option, with performance nearly identical to the Cisco 7200VXR NPE-G1. The main difference is the Cisco 7200VXR will be upgradeable at some point in the future to newer and faster processing engines and encryption modules. The Cisco 7301 is a fixed-configuration platform, and is not upgradeable.

Tunnel Aggregation Scalability

The maximum number of IPsec tunnels that a headend can terminate must be considered. Tunnel scalability is a function of the number of branch routers that are terminated to the headend aggregation point. This number needs to include both the primary tunnels as well as any alternate tunnels for which each headend may be responsible in the event of a failover situation.

The number of IPsec tunnels that can be aggregated by a platform is used as the primary determining factor in recommending a platform. Equally or more important is the encryption pps rate.

Aggregation Scalability

Aside from the number of tunnels that a headend terminates, the aggregated pps must be considered. Requirements are influenced by several factors, including the following:

- Headend connection speed—What is the speed of the WAN link on which the IPsec tunnels of the branch routers are transported through at the headend? (DS3, OC3, OC12, other?)
- Branch connection speeds—What is the typical bandwidth at each branch office going to be? (Fractional-T1, T1, T3, broadband DSL/cable, other?)
- Expected utilization—What is the maximum utilization of the WAN bandwidth under normal operation (or perhaps peak, depending on customer requirements)?

The pps rate (traffic size and traffic mix) is the largest single factor to branch router scalability.

Customer Requirement Aggregation Scalability Case Studies

This section includes examples to illustrate headend scalability factors.

Customer Example with 300–500 Branches

A customer has the design requirements shown in [Table 3-3](#):

Table 3-3 Customer Requirements

Customer Requirement	Value
Number of branch offices	300
Branch access speeds	256 Kbps FR PVCs
Headend access speed	DS3 (90 Mbps bi-directional))
Expected utilization	75%


The calculation of aggregate bandwidth requirements is as follows:

- Typical case— $300 \times 256 \text{ Kbps} \times 2 \text{ (bi-directional)} \times 80\% \text{ utilization} = 122 \text{ Mbps}$
- Worst case— $300 \times 256 \text{ Kbps} \times 2 \text{ (bi-directional)} \times 100\% \text{ utilization} = 155 \text{ Mbps}$

Even though the worst case aggregation option calculated is 155 Mbps, the total headend connection speed of DS3 (90 Mbps bi-directional) is the constraining factor. In this example, the key factor to consider is a headend router platform that can support 300 tunnels and provide an aggregate encryption bandwidth of at least 61 Mbps. The traffic mix on the network determines the pps load on the processor. [Chapter 4, “Scalability Test Results \(Unicast Only\),”](#) includes tables that use a traffic mix commonly found on enterprise networks (e-mix) to determine pps based on bps.

At a minimum, a Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator supports these requirements. For additional future growth capacity, a Cisco 7200VXR with NPE-2 processor and VSA could also be positioned. For platform-specific results, see [Headend Scalability Test Results—p2p GRE over IPsec, page 4-3](#). This design is shown in [Figure 3-1](#).

Figure 3-1 Cisco 7200VXR-Based p2p GRE over IPsec VPN Design



In this design, each Cisco 7200VXR router with NPE-G1 and SA-VAM2+ is handling half of the branch routers (250) and traffic load (45 Mbps) under normal circumstances. Both p2p GRE and IPsec are being terminated directly on each Cisco 7200VXR.

If either of the headends experiences a failure, the surviving headend can handle the total number of branches (500) and traffic load (90 Mbps) during the failure.

The headends can reside in the same or separate geographic locations. If in separate locations, it is likely that each location has independent DS3 (in this example) connections. Separate DS3s should be factored into the maximum aggregation bandwidth requirements to make sure the platforms recommended can handle the load.

Customer Example with 1000 Branches

Next, consider another customer example with many branch locations, each with relatively small traffic requirements, such as a point-of-sale terminal model, as shown in [Table 3-4](#).

Table 3-4 Customer Requirements

Customer Requirement	Value
Number of branch offices	1000
Branch access speeds	128 Kbps FR PVCs
Headend access speed	DS3 (90 Mbps bi-directional))
Expected utilization	25%

The calculation of aggregate bandwidth requirements is as follows:

- Typical case—2000 x 128 Kbps x 2 (bi-directional) x 25% utilization = 64 Mbps
- Worst case—2000 x 128 Kbps x 2 (bi-directional) x 100% utilization = 256 Mbps

In this example, the main factor to consider is a headend router platform that can support at least 1000 tunnels. Although the worst case aggregation option calculated is 256 Mbps, the aggregate headend connection speed of DS3 (90 Mbps) is the upper constraining factor. Normal utilization is expected to be in the range of 64 Mbps.

In this case, even though a Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator can support the number of tunnels required, the encryption bandwidth required exceeds performance of the platform and produces a bottleneck at approximately 90–100 Mbps.

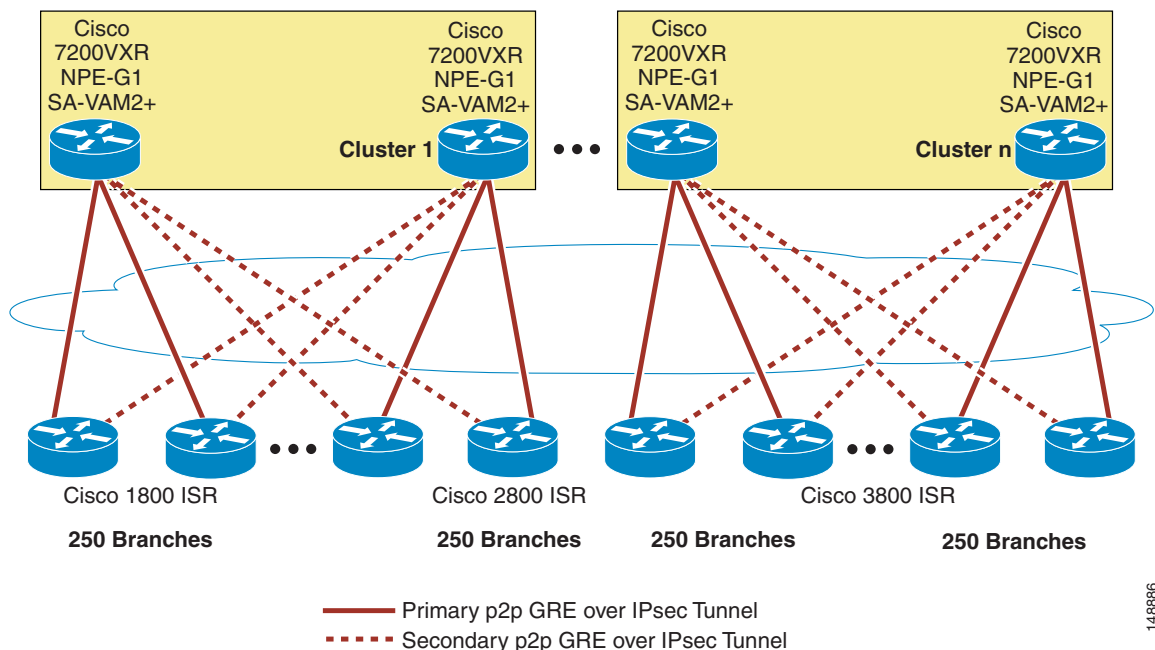
Recommendations for this customer are to do one of the following:

- Divide the tunnels aggregated across multiple Cisco 7200VXR
- Use a platform with higher encryption performance, such as the Cisco 7600 with VPN Shared Port Adapter (SPA)

A design based on the Cisco 7600 with a VPN SPA is recommended. For platform-specific results, see [Headend Scalability Test Results—p2p GRE over IPsec, page 4-3](#).

Consider the first option. The “building block” of 500 branches to a pair of Cisco 7200VXR is duplicated to support the total number of branch offices. In the customer example above, this requires two pairs of Cisco 7200VXR for a total of 1000 branches. There is no limit to the number of branches that can be aggregated with this approach. [Figure 3-2](#) shows this solution.

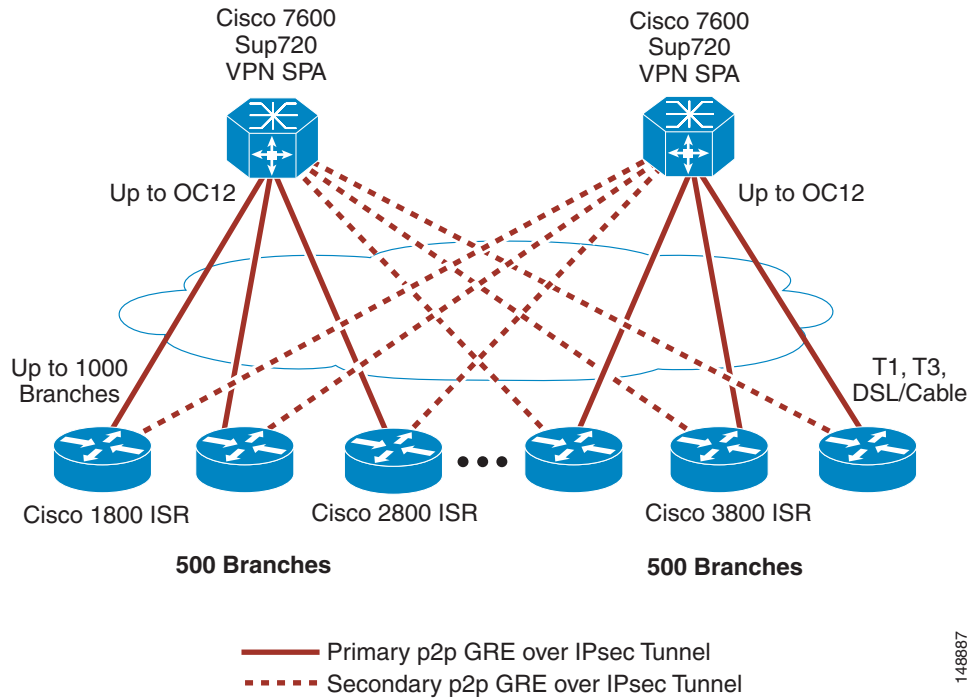
Figure 3-2 Stacked Cisco 7200VXR Design



148886

Another way to address these same customer requirements is to position a larger aggregation platform in the design. The design alternative is to implement a pair of Cisco 7600 routers, each with Sup720 and VPN SPA. Each router can support 1000 GRE tunnels. (See [Figure 3-3](#).)

Figure 3-3 Cisco 7600-Based p2p GRE over IPsec VPN Design



This design is an example of [Load Sharing with Failover Headend Resiliency Design, page 2-21](#). Each branch office has a primary p2p GRE over IPsec tunnel to one Cisco 7600, and an alternate tunnel to a second Cisco 7600. [Figure 3-3](#) can support up to 1000 branch offices, with each Cisco 7600 having 500 primary and 500 secondary neighbors. Routing neighbors tends to be the limiting factor. On the Cisco 7200VXR platform, routing neighbors tend to be limited to 500–700. On the Cisco 7600 with Sup720, up to 1000 routing neighbors have been proven to work in the Cisco scalability test lab.

The VPN SPA can hardware accelerate up to 8000 IPsec tunnels, but only 2047 p2p GRE tunnels. Encryption throughput is not a concern, because the VPN SPA can support line rates up to OC12.

Another design consideration is to determine where to terminate the p2p GRE tunnels. IPsec is terminated on the VPN SPA. The p2p GRE tunnels can either be terminated on the VPN SPA or on the Sup720. Both options were tested in the Cisco scalability lab, with the Sup720 option providing approximately 10–20 percent additional pps performance for the overall design. The design chosen depends on other factors, such as what other functions the Sup720 may be performing.

Customer Example with 1000–5000 Branches

There are cases where customers require aggregation of several thousand branch offices over a VPN. Consider the case of 5000 branch offices, each with a tunnel to a primary and alternate hub location, such as a point-of-sale terminal model. (See [Table 3-5](#).)

Table 3-5 Customer Requirements

Customer Requirement	Value
Number of branch offices	5000
Branch access speeds	128 Kbps/1 Mbps DSL
Headend access speed	OC12 (1.24 Gbps bi-directional)
Expected utilization	25%

The calculation of aggregate bandwidth requirements is as follows:

- Typical case— $5000 \times (128 \text{ Kbps} + 1 \text{ Mbps}) \times 22\% \text{ utilization} = 1.24 \text{ Gbps}$
- Worst case— $5000 \times (128 \text{ Kbps} + 1 \text{ Mbps}) \times 100\% \text{ utilization} = 5.64 \text{ Gbps}$

One design decision that needs to be made is where to terminate the p2p GRE tunnels. IPsec is terminated on the VPN SPA. The p2p GRE tunnels can either be terminated on the VPN SPA or on the Sup720. Both options were tested in the Cisco scalability lab, with the Sup720 option providing approximately 10–20 percent additional pps performance for the overall design. The design chosen depends on other factors, such as what other functions the Sup720 may be performing.

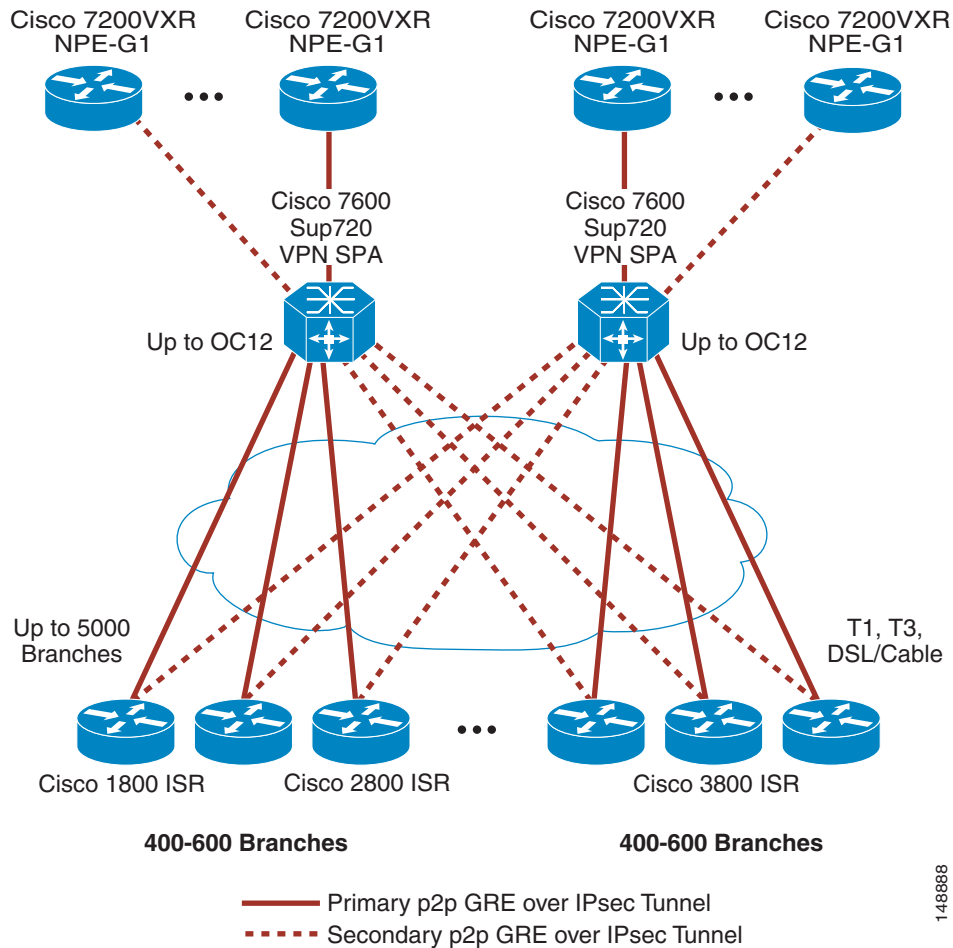
Although the worst case aggregation option calculated is 5.64 Gbps, the total headend connection speed of OC12 (1.24 Gbps bi-directional bandwidth) is the upper constraining factor; thus, oversubscribing the OC12 circuit. Normal utilization is expected to be in the range of 1.24 Gbps.

Currently, no Cisco platform can aggregate 5000 p2p GRE over IPsec tunnels on a single platform. Factors such as the routing protocol limitations cause scalability issues. 5000 IGP RP neighbors do not scale on a single platform, which contradicts the widely known accepted practices, regardless of VPN technologies. Options for such large designs include the following:

- Duplicating a smaller-scale design, such as either the Cisco 7200VXR-based design for 500 branches, or the Cisco 7600-based design for 1000 branches.
- Implementing a Dual Tier Headend Architecture, using the Cisco 7200VXR platform to terminate the p2p GRE tunnels, and the Cisco 7600 platform for high-capacity IPsec encryption.

The Dual Tier Headend Architecture is shown in Figure 3-4.

Figure 3-4 Dual Tier Headend Architecture with p2p GRE over IPsec



148888

The Cisco 7200VXR platforms terminate the p2p GRE tunnels. Because there are no IPsec encryption requirements in this “tier” of the design, no SA-VAM2+ is required, and also these platforms can typically handle more spokes than if the router was performing both p2p GRE and IPsec.

In the other “tier” of the design, the Cisco 7600 with Sup720 and VPN SPA performs IPsec encryption services, which enables a single Cisco 7600, providing up to OC12 encryption speed, to perform as the IPsec tunnel aggregation point for up to 5000 tunnels.

A very important limitation of this design approach is that IP multicast limits the total number of tunnels that can be terminated through the VPN SPA. For designs requiring IP multicast, see the *Multicast over IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Branch Office Scalability

The branch routers are primarily responsible for the following:

- Terminating p2p GRE over IPsec tunnels from the headend routers
- Running a routing protocol inside of the p2p GRE tunnels to advertise internal routes

The most important factors to consider when choosing a product for the branch office include the following:

- Branch access speed, and expected traffic throughput to the headend (Fractional T1, T1, T3, broadband cable/DSL, other?)
- What other services is the branch router providing (for example, DHCP, NAT/PAT, VoIP, Cisco IOS firewall, IOS-IPS, and so forth)?

The pps rate (traffic size and traffic mix) is the largest single factor to branch router scalability.

The number of p2p GRE over IPsec tunnels does not play a large role in the branch sizing, because each branch router must be able to terminate a single tunnel for this design topology.

A primary concern is the amount of traffic throughput (pps and bps) along with the corresponding CPU utilization. Cisco recommends that branch routers be chosen so that CPU utilization does not exceed 65 percent under normal operational conditions. The branch router must have sufficient CPU cycles to service periodic events that require processing. Examples include ISAKMP and IPsec SA establishing and re-keying, SNMP, and Syslog activities, as well as local CLI exec processing. The average CPU busy on the branch router can be higher than the crypto headend because the headend is responsible for termination of all branches, not only the requirements of the branch being serviced by the remote router.

After initial deployment and testing, it may be possible to run branch routers at CPU utilization levels higher than 65 percent under normal operational conditions. However, this design guide conservatively recommends staying at or below 65 percent.

The Cisco Integrated Services Router (ISR) 1840, 2800, and 3800 Series of products have higher CPU performance than the products they replace. The ISR has an encryption module on the motherboard, or can be upgraded to an AIM series of encryption module for increased crypto performance.



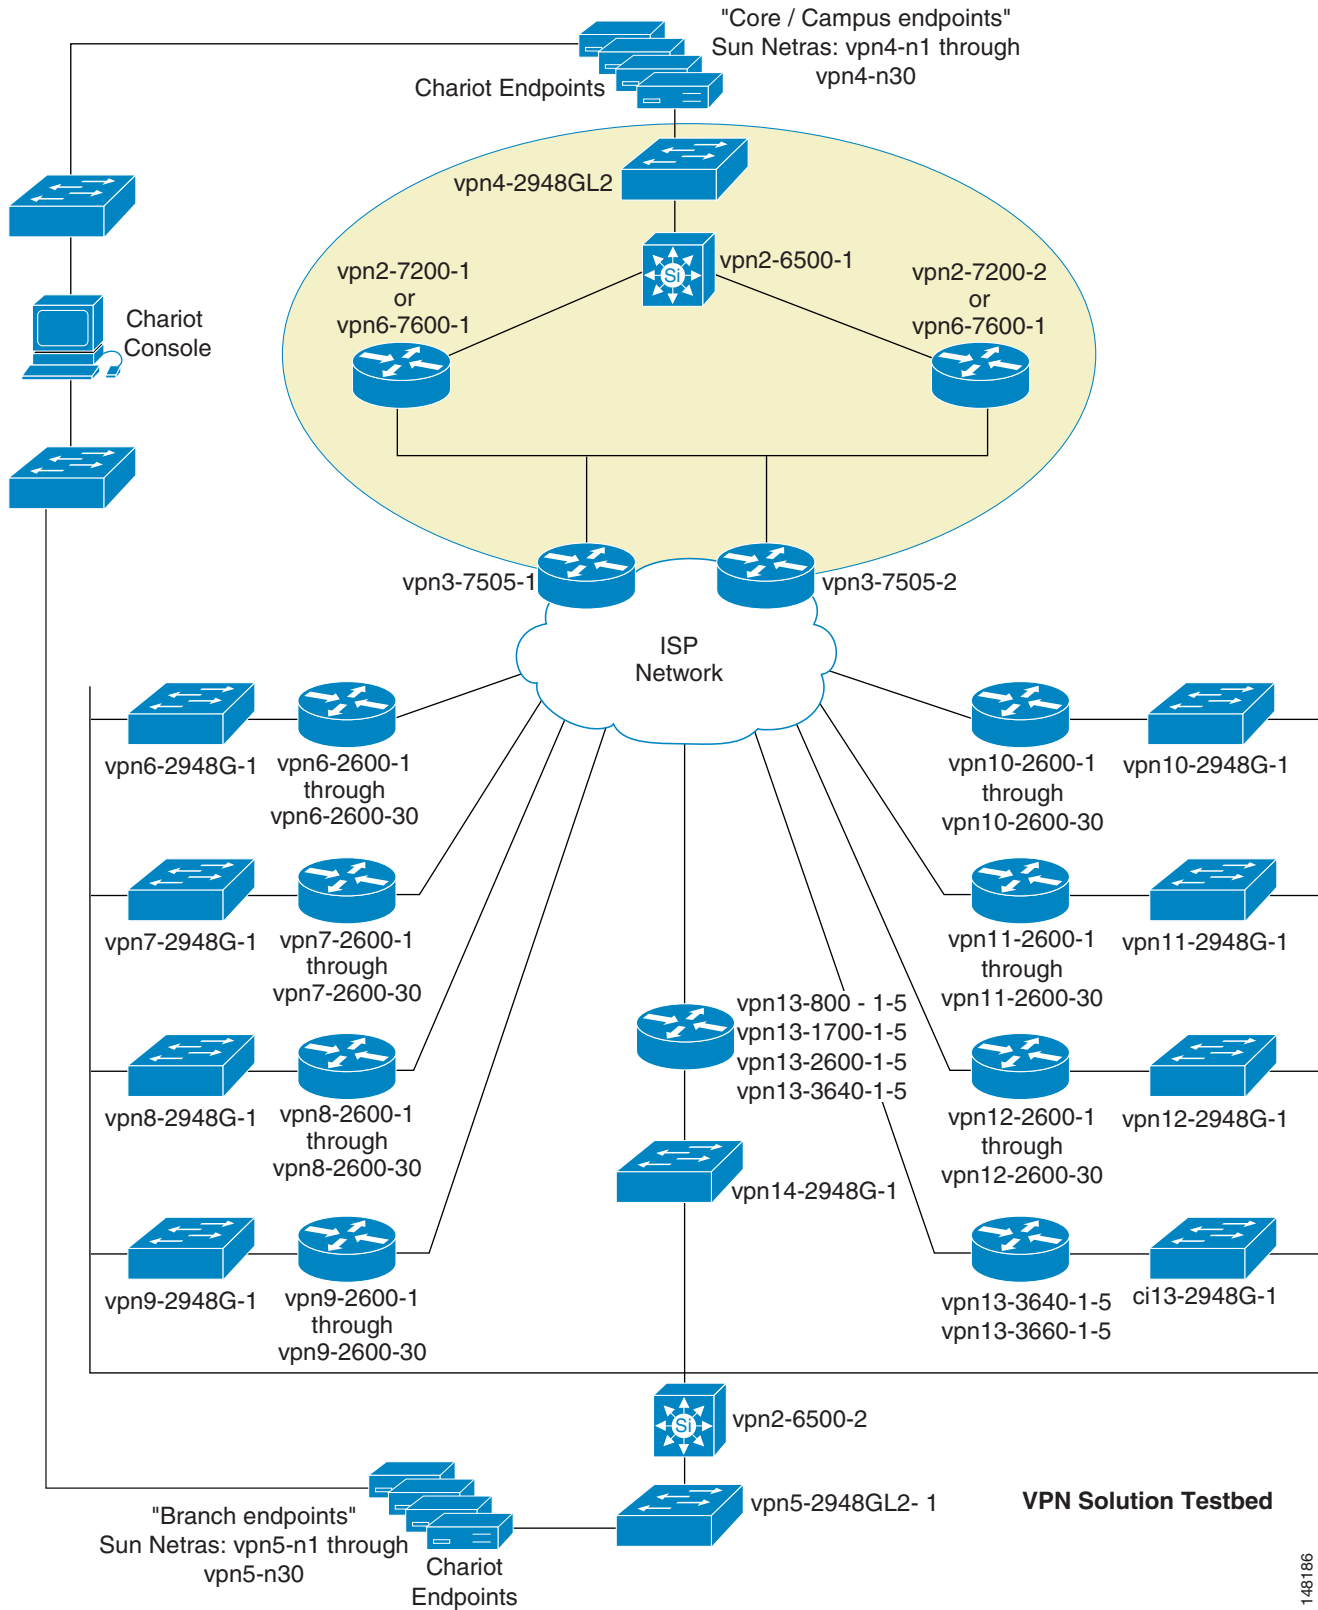
Scalability Test Results (Unicast Only)

This section offers test results from the Cisco test lab, to provide design guidance on the scalability of various platforms in p2p over IPsec VPN designs. IP multicast (IPmc) results are not included.

Scalability Test Bed Network Diagram

[Figure 4-1](#) shows the scalability test bed network diagram.

Figure 4-1 Scalability Test Bed Network



148186

Scalability Test Methodology

The headend scalability test bed consists of a number of Cisco branch routers (various types, including the 1700, 2600, 3600, 3700, 1800, 2800, and 3800 families) homed to various types of headends. For most of the traffic sent through the network, flows are established using the Ixia Chariot testing tool. The bps mix of traffic is approximately 35 percent UDP and 65 percent TCP; application types represented in the mix include the following: VoIP, FTP, DNS, HTTP, POP3, and TN3270. The average packet size is 188 bytes, from headend to branch, and 144 bytes from branch to headend. This relatively small average packet size ensures that the scalability results presented support a converged network design, and tends to be fairly conservative. A network carrying data-only traffic, with a larger average packet size, may achieve better bps performance than that listed here. However, the pps performance given a specific CPU value should be the same.

Some traffic is also generated by the Cisco IP SLA feature in IOS, formerly known as Cisco Service Assurance Agent (SAA), using the HTTP Get script, with the branch routers making an HTTP Get call to an HTTP server in the core. Testing was conducted without fragmentation occurring in the network by setting the MTU to 1300 bytes on the test endpoints.

The following tables show results for testing with a configuration for p2p GRE over IPsec tunnel aggregation. The routing protocol used during testing was EIGRP unless otherwise stated. The traffic mix used, as stated earlier, is converged data and g.729 VoIP.

Headend Scalability Test Results—p2p GRE over IPsec

Table 4-1 shows results for the testing with the configuration for a p2p GRE over IPsec design and no other Cisco IOS features such as IOS Firewall, PAT, ACLs, IPS, or QoS.

Table 4-1 Scalability Test Results—p2p GRE over IPsec Only

Platform	# of Tunnels	# Voice Calls	Throughput (Kpps)	Throughput (Mbps)	CPU %
Cisco 7200VXR with NPE-G1 and dual SA-VAM2	500	240	40.0	82.5	80%
Cisco 7600 Sup720 VPN SPA	1000 (Both p2p GRE tunnels and IPsec tunnels on VPN SPA)	3537	480.0	1050.0	N/A
	1000 (p2p GRE tunnels on Sup720 and IPsec tunnels on VPN SPA)	4137	521.0	1110.0	N/A
Cisco 7200VXR and Cisco 7600 Dual Tier Headend Architecture	3000 (1000 p2p GRE tunnels on each of three Cisco 7200VXR with IPsec tunnels on VPN SPA)	est. 4000	601 in total Up to 203 Kpps on each of three 7200VXR	-	N/A

Note that headend scalability testing did not include an exhaustive evaluation of the maximum number of tunnels that can be terminated to headend devices. In addition, scalability testing of the branch routers was performed with two tunnels per branch. This did not include exhaustive testing of the number of tunnels these various platforms can support.

Headend Scalability Test Results—p2p GRE Only

Table 4-2 shows results for the testing with the configuration for a p2p GRE only design and no other Cisco IOS features such as IOS firewall, PAT, ACLs, IPS, or QoS. The purpose for these results is to provide the reader with the p2p GRE only results for designing a Dual Tier Headend Architecture. The IPsec only results can be found in the *IPsec Direct Encapsulation Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Table 4-2 Scalability Test Results—p2p GRE Only

Platform	# of Tunnels	# Voice Calls	Throughput (Kpps)	Throughput (Mbps)	CPU %
Cisco 7200VXR with NPE-G1	1000	1377	203.7	414.2	79%
Cisco 7600 Sup720 ¹	1000	4437	665.7	1675.3	N/A

1. These results are limited by the test bed traffic load limitation. This platform is expected to exceed these results.

Branch Office Scalability Test Results

Table 4-3 shows results for testing with a configuration for p2p GRE over IPsec. A single tunnel was configured to the aggregation headend. These results include other integrated Cisco features such as IOS Firewall, PAT, and ACLs, but not QoS or IPS.

Table 4-3 Branch Office Scalability Test Results


Platform	HW Encryption	# Voice Calls	Throughput (Kpps)	Throughput (Mbps)	CPU %
Cisco 3845 ISR	On-board	187	24.0	48.8	81%
	AIM-VPN/HPH-Plus	420	27.1	50.1	80%
Cisco 3825 ISR	On-board	143	18.2	36.6	81%
	AIM-VPN/EH-Plus	156	20.1	42.8	79%
Cisco 2851 ISR	On-board	90	11.4	23.8	79%
	AIM-VPN/EH-Plus	120	14.9	30.8	80%
Cisco 2821 ISR	On-board	45	6.0	13.6	53%
	AIM-VPN/EH-Plus	97	12.3	25.9	78%
Cisco 2811 ISR	On-board	19	2.6	5.8	79%
	AIM-VPN/EH-Plus	27	3.6	8.0	80%

Table 4-3 Branch Office Scalability Test Results

Cisco 2801 ISR	On-board	19	2.6	5.8	83%
	AIM-VPN/EPII-Plus	30	3.9	8.4	79%
Cisco1841 ISR	On-board	19	2.5	5.7	82%
	AIM-VPN/BPII-Plus	30	3.9	8.8	80%
Cisco 1811W with no BVI configured	On-board	33	7.6	16.0	81%
Cisco 1811W with BVI configured	On-board	60	4.3	9.3	82%
Cisco 871W with no BVI configured	On-board	8	2.0	4.4	85%
Cisco 871W with BVI configured	On-board	15	1.1	2.4	84%

AES versus 3DES Scalability Test Results

Both 3DES and AES encryption are available in all products shown here, including hardware-accelerated IPsec. Not every test was executed with both 3DES and AES; however, several snapshot tests were performed to compare performance. As shown in [Figure 4-2](#), results are relatively comparable, with little to no variation in performance, even for AES with wider key lengths.

Figure 4-2 Comparison of 3DES and AES Performance

Failover and Convergence Performance

Customers may have different convergence time requirements. The design principles in this guide were used to perform scalability tests with up to 500 branch offices aggregated to two or three headend routers. See [Load Sharing with Failover Headend Resiliency Design, page 2-21](#) for more details.

The following test was performed by powering off one of the headend devices to simulate a complete headend failure. The network fully converged after a maximum of approximately 32 seconds for all 240 branches. The test scenario is shown in [Table 4-4](#).

Table 4-4 Three Headend Load Sharing with Failover

	Headend 1	Headend 2	Headend 3
Cisco 7200VXR NPE-400	Cisco IOS version 12.1(9)E	Cisco IOS version 12.1(9)E	Cisco IOS version 12.1(9)E
Starting condition	27 Mbps 80 branches 32% CPU	28 Mbps 80 branches 32% CPU	44 Mbps 80 branches 37% CPU
During failover	Powered off	41 Mbps 120 branches 46% CPU	56 Mbps 120 branches 50% CPU

The same test was then performed with 500 branch offices aggregated to two headend devices. All 250 branches from the powered-off headend successfully failed over to the single surviving headend. In this test, the network fully converged after approximately 32 seconds.

During the re-convergence, when the powered-off headend was restored, the convergence of each branch took approximately two seconds each, with the total time for re-convergence at about five and one-half minutes. [Table 4-5](#) shows the resulting CPU utilization percentages.

Table 4-5 Two Headend Load Sharing with Failover

	Headend 1	Headend 2
Cisco 7200VXR NPE-G1	Cisco IOS version 12.2(13)S	Cisco IOS version 12.2(13)S
Starting condition	37.5 Mbps 250 branches 43% CPU	35.6 Mbps 250 branches 43% CPU
During failover	Powered off	45.7 Mbps 500 branches 84% CPU

Because of the normal TCP backoff process, the total traffic levels through the surviving router may be temporarily lower after a failure than the total traffic before failure.

Software Releases Evaluated

Table 4-6 shows the software releases used in the scalability testing.

Table 4-6 **Software Releases Evaluated**

Cisco Product Family	SW Release
Cisco 7600 VPN SPA	Cisco IOS 12.2(18)SXE2
Cisco Catalyst 6500 VPNSM	Cisco IOS 12.2(17d)SXB1
Cisco 7200VXR	Cisco IOS 12.2(13)S Cisco IOS 12.1(9)E Cisco IOS 12.3(5)
Cisco branch office routers (17xx, 26xx, 36xx, 37xx)	Cisco IOS 12.3(8)T5
Cisco branch office ISRs (1841, 28xx, 38xx)	Cisco IOS 12.3(8)T5 Cisco IOS 12.3(11)T2
Cisco remote office routers (831)	Cisco IOS 12.3(8)T5
Cisco remote office routers (871W and 1811W)	Cisco IOS 12.3(14)YT1

Before selecting Cisco IOS software, perform the appropriate research on cisco.com, and if you have technical questions, consult with Cisco Customer Advocacy (TAC).



Case Studies

The following two case studies are provided as reference material for implementing p2p GRE over IPsec designs.

Static p2p GRE over IPsec with a Branch Dynamic Public IP Address Case Study

This case study explains how to encrypt p2p GRE tunnels with a dynamic crypto map to support deployments where the remote IP address is not statically defined. This technique can also be used when there is static IP addressing on the branch office where the goal of the network administrator is to simplify the headend configuration.

Overview

To fully understand the mechanics of this configuration, it is helpful to visualize the topology as a tunnel within a tunnel. The IP addresses of the crypto endpoints are routable over the network. The Internet is the most common deployment. The headend crypto IP address is a static IP address. The IP address of the outside interface of the branch router is also routable over the Internet but is dynamically assigned by the broadband or Internet service provider, either through DHCP or PPPoE. The branch router always initiates the p2p GRE and crypto tunnel to the crypto headend IP address. The interesting traffic (traffic matching the crypto ACL) is either a GRE keepalive or hello packet of whatever routing protocol is configured on the branch router.

The crypto headend router learns the dynamically assigned IP address of the outside interface of the branch router because it is the source IP address in the ISAKMP packets. Following successful negotiation of IKE Phase 1 (main mode) and subsequently IKE Phase 2 (quick mode), the outer (crypto) tunnel is established providing transport for the p2p GRE tunnel.

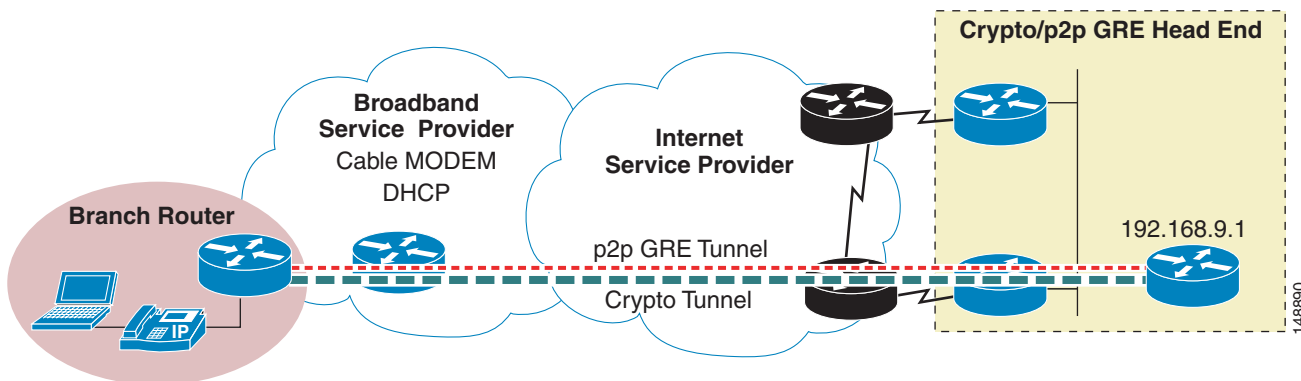
The p2p GRE tunnel is transported within the crypto tunnel. The p2p GRE tunnel endpoint IP addresses are only locally significant. The tunnel source is a connected interface, and the tunnel destination IP address must be routed out the interface with the crypto map configured. After the encrypting router has identified the p2p GRE packet as requiring encryption, and is encapsulated in IPsec, the packet is routed on destination IP address in the IPsec header and not the destination address in the p2p GRE header.

As the encrypted packet is forwarded over the network, the p2p GRE destination IP address is simply part of the IPsec payload. All routing decisions are based on the IPsec header. When the IPsec packet is decrypted by the receiving router, the packet forwarding decision is then based on the p2p GRE destination IP address. This p2p GRE destination address is a connected interface on the decrypting router.

Sample Topology

A typical customer topology includes two p2p GRE tunnels from the branch router to separate crypto headend routers. One tunnel is the primary tunnel and the second is the backup tunnel providing redundancy in the event the primary headend or network connection failure. To simplify the explanation, only one branch router, headend, and tunnel is shown in this section and in [Figure 5-1](#).

Figure 5-1 Sample Topology



Addressing and Naming Conventions

In this section, assume that 10.0.0.0/8 address space is the enterprise IP address space and 192.168.0.0/16 represents routable Internet address space.

The addressing scheme and naming conventions are defined as follows:

- Branch
 - Routers in this example are assigned address space within a /22 prefix.
 - This address space is advertised as a summary route to the headend(s).
 - Loopback 0 is used for p2p GRE tunnel source IP address.
 - Loopback 1 is used for Network Management.
 - Tunnel 1 is the p2p GRE tunnel to the primary headend.
 - Tunnel 0 is the p2p GRE tunnel to the secondary headend (if present).
 - The last octet of the IP address of Loopback 0 corresponds to the branch site number (17,18, and so on).

- Headend (primary)
 - Loopback 0 IP address is Internet routable represented by 192.168.9.1.
 - The p2p GRE tunnel interface instance number represents the branch router supported by this interface. For example, branch_17 is Tunnel17, and branch_18 is Tunnel18, and so on.
 - A static or dynamically learned route for 10.63.0.0/16 is present to switch packets out the interface with the dynamic crypto map configured.
- Headend (secondary—configuration not shown)
 - Loopback 0 IP address is Internet routable represented by 192.168.8.1.
 - *All primary headend assumptions also apply to this secondary headend.*

It is a best practice for a network manager to document and define an address plan before implementation. Table 5-1 shows an example. The configuration for Branch 17 is also shown in the configuration details. The data for Branch 18 is shown to illustrate how the addressing scheme progresses for subsequent branch routers. The configuration for Branch 18 or the secondary headend router is not shown.

Table 5-1 Address Plan

					Primary	Headend
Branch 17		Loopback0	10.63.0.17/32	10.63.0.17	192.168.9.1	Loop0
	10.0.4.0/22	Branch network	Summary	?		
		FastEth0/0	10.0.4.0/24			
		Loopback1	10.0.6.1/32	10.0.6.1		
		Tunnel 1	10.0.7.4/30	10.0.7.5	10.0.7.6	Tunnel 17
		Tunnel 0	10.0.7.0/30	10.0.7.1	10.0.7.2	
Branch 18		Loopback0	10.63.0.18/32	10.63.0.18	192.168.9.1	Loop0
	10.0.8.0/22	Branch network	Summary	?		
		FastEth0/0	10.0.8.0/24			
		Loopback1	10.0.10.1/32	10.0.10.1		
		Tunnel 1	10.0.11.4/30	10.0.11.5	10.0.11.6	Tunnel 18
		Tunnel 0	10.0.11.0/30	10.0.11.1	10.0.11.2	

Configuration Examples

The following sections provide configuration examples.

p2p GRE Tunnel and Interface Addressing

Based on [Table 5-1](#), a partial configuration example for the branch router (branch_17) is shown.

Items to note include the following:

- The branch router has a static route for 192.168.8.1/32 and 192.168.9.1/32 hosts with the DHCP learned default gateway as the next hop. This network contains the IP address of both the primary and secondary headend addresses 192.168.8.1 and 192.168.9.1.
- A default route is learned through the dynamic routing protocol running in the p2p GRE tunnel following crypto tunnel establishment and routing protocol neighbors are formed on the p2p GRE tunnel and routes are exchanged. The 192.168.8.1/32 and 192.168.9.1/32 routes are a best practice to avoid recursive routing.

```

!
hostname branch_17
!
interface Loopback0
  description For Tunnel Termination
  ip address 10.63.0.17 255.255.255.255
!
interface Loopback1
  description For Network Management
  ip address 10.0.6.1 255.255.255.255
!
interface FastEthernet 0/0
  description Inside (LAN) interface
  ip address 10.0.4.1 255.255.255.0
!
interface Ethernet1/0
  description Outside to CABLE MODEM / Broadband SP / Internet
  ip address dhcp
  crypto map GRE
!
interface Tunnell
  description Tunnel to (Primary Headend)
  ip address 10.0.7.5 255.255.255.252
  ip summary-address eigrp 44 10.0.4.0 255.255.252.0 5
  tunnel source Loopback0
  tunnel destination 192.168.9.1
!
router eigrp 44
  network 10.0.4.0 0.0.3.255
!
ip route 192.168.8.1 255.255.255.255 dhcp
ip route 192.168.9.1 255.255.255.255 dhcp
!
end

```

The primary headend outside tunnel (for branch 17) and loopback interfaces are shown. The inside interface as well as tunnel interfaces for other branches are not shown.

Note that both the primary and secondary headends are configured using EIGRP as the IGP; however, to emphasize the routing requirements, two static routes are shown. The first is a default (0/0) route using the enterprise campus Internet gateway as a next hop. The second route is a route with a /16 prefix for network 10.63.0.0, also to the enterprise campus Internet gateway as a next hop. This route is not

required, because the default (0/0) route accomplishes that same purpose. It is included to draw attention to the fact the tunnel source IP addresses of the branch routers are allocated from the 10.63.0.0 address space as individual /32 networks. This /16 prefix route illustrates that the encrypting router selects the p2p GRE packets for encryption based on the fact the p2p GRE encapsulated packet is routed out the interface (outside) that contains the crypto map configuration.

```
hostname primary_headend
!
interface Loopback0
 ip address 192.168.9.1 255.255.255.0
!
interface Ethernet1/0
 description Outside Interface
 ip address 10.254.1.49 255.255.255.0
 crypto map DYNO-MAP
!
interface Tunnel17
 ip address 10.0.7.6 255.255.255.252
 tunnel source Loopback0
 tunnel destination 10.63.0.17
!
ip route 0.0.0.0 0.0.0.0 10.254.1.1 name To_INTERNET_GateWay
ip route 10.63.0.0 255.255.0.0 10.254.1.1 name Branch_Tunnel_Endpoints
end
```

Crypto Map Configurations (Crypto Tunnel)

The crypto map of the branch router is simply a static crypto map matching the p2p GRE tunnel endpoint IP addresses. Only the primary headend is shown.

```
!
hostname branch_17
!
interface Loopback0
 ip address 10.63.0.17 255.255.255.255
!
crypto map GRE 10 ipsec-isakmp
 set peer 192.168.9.1
 set transform-set AES_SHA_TUNNEL 3DES_SHA_TUNNEL
 match address GRE_to_NINE
!
ip access-list extended GRE_to_NINE
 permit gre host 10.63.0.17 host 192.168.9.1
!
```

The primary headend is configured with a dynamic crypto map. Note there is no **match address** configured. It is learned dynamically from the branch router.

The **match address** ACL on the crypto headend crypto map is optional. If a remote site has not established its crypto tunnel, the headend router sends p2p GRE packets out the crypto map interface unencrypted. However, this is a minimal security exposure for two reasons: the destination IP address of the unencrypted p2p GRE packet in this example is an RFC1918 IP address and is not routed over the Internet by the ISP; and the contents of these packets are either GRE keepalives or routing protocol hello packets. Data packets are not routed until the p2p GRE tunnel is up and the routing protocol has formed a neighbor relationship.

```
hostname primary_headend
!
crypto dynamic-map DYNO-TEMPLATE 10
 description dynamic crypto map
```

```

set transform-set AES_SHA_TUNNEL 3DES_SHA_TUNNEL
!
crypto map DYNO-MAP local-address Loopback0
crypto map DYNO-MAP 10 ipsec-isakmp dynamic DYNO-TEMPLATE

```

Note that crypto maps are not required on tunnel interface beginning in 12.2(13)T.

Headend EIGRP Configuration

The headend routers determine what networks are learned by the branch router through the tunnel. This is effectively a policy push from the headend to all branch routers. The branch router learns a default route if the outside IP address is obtained by DHCP, and this default route is inserted in the routing table of the branch router by default with an administrative distance of 254 or least preferred.

The remote router is configured to encrypt p2p GRE packets. The networks advertised by the headend through the tunnel interface determine what gets encrypted or what is routed to the default route to the Internet. If the headend router advertises a default (0/0) route through the p2p GRE tunnel, split-tunnel is not implemented on the branch router. If the headend router advertises, for example, 10.0.0.0/8, split-tunnel is enabled and all packets not destined for 10.0.0.0/8 are routed to the Internet. Of course, the network manager should include an inbound access list as one component in securing the branch router, and this along with Context-Based Access Control (CBAC, the Cisco IOS firewall feature) controls if the return packets are permitted into the branch network.

The following is an example of the headend configuration advertising only a default route to branch 17.

```

hostname primary_headend
!
router eigrp 44
 network 10.0.0.0
 distribute-list Quad_ZERO_to_BRANCH out Tunnel17
!
ip access-list standard Quad_ZERO_to_BRANCH
 permit 0.0.0.0
!

```

It is assumed the network manager has either a default route learned from an EIGRP AS 44 neighbor or the static route to the default network is redistributed into this configuration.

Verification

The branch_17 router has an outside IP address of 192.168.31.2. Assume this is an Internet routable IP address and was learned by DHCP.

As an illustration, the primary headend dynamic crypto map is shown:

```

primary_headend#show crypto map
Crypto Map: "DYNO-MAP" idb: Loopback0 local address: 192.168.9.1

Crypto Map "DYNO-MAP" 10 ipsec-isakmp
    Dynamic map template tag: DYNO-TEMPLATE

Crypto Map "DYNO-MAP" 65536 ipsec-isakmp
    Peer = 192.168.31.2
    Extended IP access list
        access-list permit gre host 192.168.9.1 host 10.63.0.17
    dynamic (created from dynamic map DYNO-TEMPLATE/10)
    Current peer: 192.168.31.2

```

Note in the above display that the dynamic crypto map entry has been updated with the crypto map ACL (source and destination IP address are reversed accordingly) configured on branch_17.

The ISAKMP SA is displayed for branch_17, showing that the source IP address of the ISAKMP session is the address learned by DHCP.

```
branch_17# show crypto isakmp sa
dst          src          state        conn-id slot
192.168.9.1  192.168.31.2 QM_IDLE      37      0
```

Summary

Using the configuration guide in this section, network managers can implement p2p GRE tunnels in network topologies where the IP address of the branch router is learned dynamically. Static crypto maps are configured on the headend router; as branch routers are brought online, no change to the headend crypto configuration is required. The headend router needs only a corresponding p2p GRE tunnel interface for the new branch.

The primary advantage for this configuration compared to a DMVPN configuration is a p2p GRE interface as opposed to an mGRE interface. P2p GRE interfaces support multiprotocol (IPX, Appletalk, and so on) routing as well as interface (and peer) specific configurations such as a QoS service policy.

Moose Widgets Case Study

The key objective of this case study is to provide a reference example for a site-to-site VPN design. It provides an example of how these design principles can be applied in a real-world scenario.

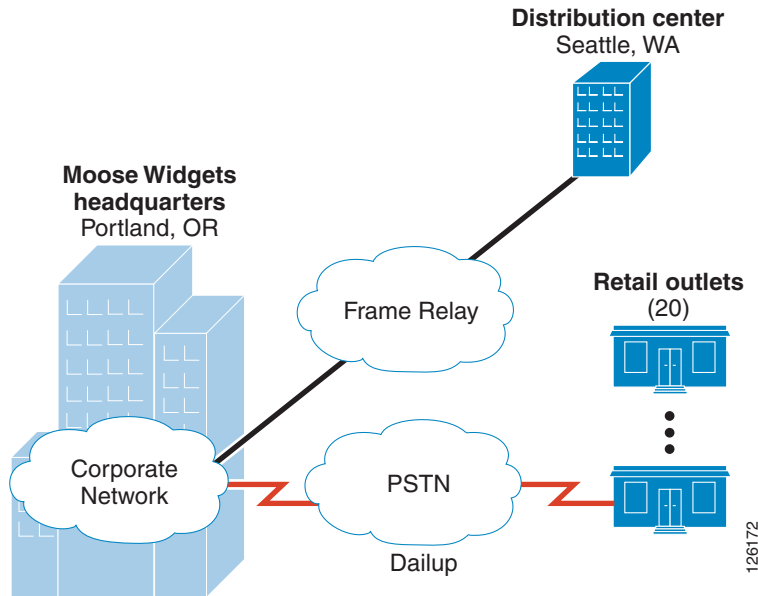
The details of the service provider backbone and WAN connectivity are not addressed in the case study, because the focus is on VPN deployment on the enterprise customer side.

Customer Overview

Moose Widgets has been developing products at their Portland, Oregon headquarters (HQ) for several years. In addition, Moose has a single distribution center and 20 retail outlets across the United States (US).

Currently, Moose Widgets uses a traditional Frame Relay (FR) WAN service to connect its headquarters to its distribution center. There is currently no connectivity to retail outlets, with the exception of a few outlets that use a personal computer (PC) to dial-up to the corporate HQ. The current network topology is shown in [Figure 5-2](#).

Figure 5-2 Moose Widgets Case Study—Current Topology



Moose has recently acquired two companies; one in San Jose, California and the other in Great Falls, Montana. Moose wants to connect its newly acquired companies and its retail outlets to its corporate network as an intranet. In addition, Moose plans to expand its retail outlets to 40–50 over the next year and sees already that it will most likely need additional distribution centers on the East and West coasts of the US.

As part of a corporate initiative, Moose is implementing a centralized inventory tracking system to better manage inventory in its growing distribution centers and retail outlets to significantly lower costs. The existing dial-in access does not provide adequate bandwidth to support the new applications. Further, Moose is concerned about escalating dial-in charges as each retail outlet relies more on corporate resources. As a result, Moose is looking to transition to a dedicated connection for each of its retail outlets using the Internet and VPN technology.

Moose is concerned about the costs of adding the connections. They are also concerned about the ability to quickly get retail outlets up and running. Moose indicates they are primarily concerned about data traffic today, but there is some degree of interest in adding voice services in the future.

Moose estimates their traffic requirements for their various site locations as shown in Table 5-2.

Table 5-2 Moose Widgets Case Study—Bandwidth

Location	Connectivity Requirements
Distribution center (one today, potentially three in the future)	T1 (3 Mbps bi-directional)
San Jose	2 x T1 (6 Mbps bi-directional)
Great Falls	2 x T1 (6 Mbps bi-directional)
Retail outlets (up to 50)	384 Kbps/1.5 Mbps broadband DSL and cable

Design Considerations

A site-to-site IPsec VPN will be deployed with the Moose corporate HQ serving as the headend, and all other locations treated as branch sites. This allows a branch office to subscribe to a local ISP, get authenticated, and be inside the corporate intranet.

At the same time, end-to-end encryption is attained using IPsec tunneling. Switching to VPN offers Moose significant cost savings over dial-up solutions and the ability to outsource to a service provider who has VPN service as a core competency, providing more efficiency with cost and scalability.

Following the design practice outlined in [Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation,”](#) and [Chapter 3, “Scalability Considerations,”](#) there are four main design steps to perform:

- Preliminary design considerations assessing customer requirements and expectations
- Sizing of the headend devices and product selection
- Sizing of the branch site devices and product selection
- Tunnel aggregation and load distribution planning

Preliminary Design Considerations

The design is straightforward and offers flexibility. As new retail locations are put into service, Moose can purchase Internet connectivity from the local ISP, deploy a Cisco VPN router at the branch site, configure the IPsec tunnels to the headend devices at the corporate headquarters, and be up and running in a short amount of time.

Using the questions from [Design Considerations, page 2-1](#), [Table 5-3](#) summarizes the preliminary design considerations.

Table 5-3 Preliminary Design Considerations

Question	Answer	Comments
What applications does the customer expect to run over the VPN?	Data	Interested in future voice services
Is multiprotocol support required?	Yes, IP and IP multicast	GRE tunnels will enable multi-protocol traffic transport.
How much packet fragmentation does the customer expect on their network?	Minimal	Path MTU discovery enabled
How many branches does the customer expect to aggregate to the headend?	55 sites	
What is customer expected traffic throughput to/from branch offices?	See Table 5-2	
What are customer expectations for resiliency?	Resiliency is required	1 primary, 1 backup tunnel
What encryption level is required?	3DES	

Table 5-3 Preliminary Design Considerations (continued)

What type of IKE authentication method will be used?	The use of pre-shared keys is selected because of the relatively small number of sites to manage.	Migration to digital certificates should be considered if number of branches increases beyond 50 in the future.
What other services will be run on the branch VPN routers?	None	

EIGRP is recommended as the routing protocol, with route summarization.

Sizing the Headend

Although the traffic loads involved do not exceed the recommended capacity of a single headend device, Moose has indicated they would like redundancy built-in at the central location. The tunnels from the remote ends will be allocated to each of the headend devices to balance the traffic load. Secondary tunnels will also be configured and allocated so that, in the event of a headend failure, traffic will be transitioned over to the partner headend device.

Applying the sizing algorithm defined in section [Headend Scalability, page 3-3](#), the calculation of headend sizing based on number of GRE tunnels is as follows:

$$N = 55$$

$$T = N \times 2 = 110$$

$$C(t) = (T / 500) \text{ rounded up} + 1 = 110/500 \text{ rounded up} + 1 = 1 + 1 = 2 \text{ headends}$$

Next, using the throughput estimates from [Table 5-4](#), the calculation of headend sizing based on branch traffic throughput is as follows:

$$A = (3 \times 3 \text{ Mbps}) + 6 \text{ Mbps} + 6 \text{ Mbps} + (50 \times 1.9 \text{ Mbps}) = 115 \text{ Mbps} \times 50\% \text{ utilization} = 58 \text{ Mbps}$$

$$H = 109 \text{ Mbps (for Cisco 7206VXR NPE-G1 with SA-VAM2)}$$

$$C(a) = A/H, \text{ rounded up} + 1 = 58/109 \text{ rounded up} + 1 = 1 + 1 = 2 \text{ headends}$$

Comparing the number of headend devices calculated based on number of tunnels, $C(t)$, to the number based on aggregate throughput, $C(a)$, the outcomes match. Therefore it is appropriate to deploy two headend devices.

Presented with the headend product options, the customer selects to deploy two Cisco 7206VXR NPE-G1s, each equipped with an SA-VAM2 hardware encryption adapter.

Sizing the Branch Sites

The primary consideration for sizing of branch office sites is expected traffic throughput. Accordingly, starting with [Table 5-2](#), and applying the concepts presented in [Branch Office Scalability, page 3-9](#), the branch products selected are summarized in [Table 5-4](#).

Table 5-4 Moose Widgets Case Study—Branch Devices

Location	Estimated Throughput	Branch Office Platform Selected
Distribution centers	3 Mbps	Cisco 2851 ISR

Table 5-4 *Moose Widgets Case Study—Branch Devices*

San Jose	6 Mbps	Cisco 3845 ISR
Great Falls	6 Mbps	Cisco 3845 ISR
Retail outlets (typical)	384 Kbps/1.5 Mbps DSL/cable	Cisco 2801/1841 ISR

At each of the acquired company locations, a Cisco 3845 ISR is deployed. The choice of the Cisco 3845 platform is based on the assumption that the acquired companies are large offices with a substantial number of employees. Future VoIP expansion is also a factor.

At each of the distribution centers, a Cisco 2851 ISR is deployed. Finally, at each of the retail locations, a combination of Cisco 2801 ISR and 1841 ISR are deployed, depending on the size of the retail outlets.

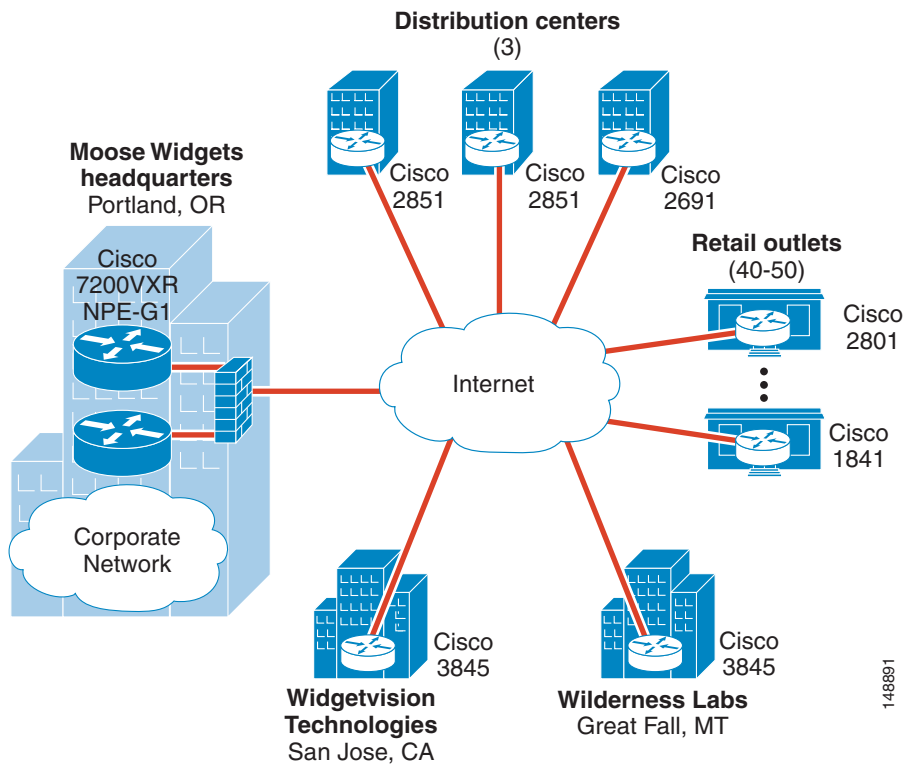
Tunnel Aggregation and Load Distribution

Given 55 branch sites, the total number of tunnels that need to be aggregated is 110 (primary and secondary). Therefore, the first headend device is allocated 27 primary and 28 backup tunnels, while the second headend device is allocated 28 primary and 27 backup tunnels.

Network Layout

The new network topology is shown in [Figure 5-3](#).

Figure 5-3 *Moose Widgets Case Study—VPN Topology*





Scalability Test Bed Configuration Files

The configurations for the central and branch sites are listed in the following sections. Note that these configurations have been extracted from real configurations used in Cisco scalability testing. They are provided as a reference only.

Cisco 7200VXR Headend Configuration

There are two headend devices in the test bed, each terminating a p2p GRE over IPsec tunnel from all branch routers. The configuration shown below is an excerpt of the first headend and does not contain configuration commands for all branches. The ISAKMP PSK, the crypto peer, the tunnel interface, and the crypto access list are shown for one device.

Headend #1:

```
ip cef
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.0.2
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto map static-map local-address GigabitEthernet0/1
crypto map static-map 100 ipsec-isakmp
  set peer 192.168.0.2
  set transform-set vpn-test
  match address b000
!
interface Loopback0
  description Loopback0
  ip address 10.57.1.255 255.255.255.255
!
interface Tunnel0
  description vpn5-2600-1-000
  bandwidth 1536
  ip address 10.60.0.193 255.255.255.252
  ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
  tunnel source 192.168.251.1
  tunnel destination 192.168.0.2
  crypto map static-map
!
interface GigabitEthernet0/1
  description GigabitEthernet0/1
```

```

ip address 192.168.251.1 255.255.255.248
duplex auto
speed auto
media-type gbic
negotiation auto
crypto map static-map
!
interface GigabitEthernet0/2
description GigabitEthernet0/2
ip address 10.57.1.1 255.255.255.248
duplex auto
speed auto
media-type gbic
negotiation auto
!
router eigrp 1
network 10.0.0.0
no auto-summary
!
ip route 0.0.0.0 0.0.0.0 192.168.251.2
!
ip access-list extended b000
permit gre host 192.168.251.1 host 192.168.0.2
!

```

Cisco Catalyst 6500/Sup2/VPNSM Headend Configuration

Headend #1:

```

hostname vpn4-6500-2
!
logging snmp-authfail
logging buffered 65535 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
clock calendar-valid
ip subnet-zero
!
no ip domain-lookup
!
mpls ldp logging neighbor-changes
mls flow ip destination
mls flow ipx destination
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.0.2
crypto isakmp key bigsecret address 192.168.1.2
crypto isakmp key bigsecret address 192.168.2.2
!
! . . . repetitive lines omitted . . .
!
crypto isakmp key bigsecret address 192.168.60.26
crypto isakmp key bigsecret address 192.168.61.26
crypto isakmp key bigsecret address 192.168.62.26
crypto isakmp keepalive 10
!

```

```
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto map static-map local-address Vlan100
crypto map static-map 100 ipsec-isakmp
  set peer 192.168.0.2
  set transform-set vpn-test
  match address b0000
crypto map static-map 101 ipsec-isakmp
  set peer 192.168.1.2
  set transform-set vpn-test
  match address b0001
crypto map static-map 102 ipsec-isakmp
  set peer 192.168.2.2
  set transform-set vpn-test
  match address b0002
!
!. . . repetitive lines omitted . . .
!
crypto map static-map 1120 ipsec-isakmp
  set peer 192.168.60.26
  set transform-set vpn-test
  match address b1020
crypto map static-map 1121 ipsec-isakmp
  set peer 192.168.61.26
  set transform-set vpn-test
  match address b1021
crypto map static-map 1122 ipsec-isakmp
  set peer 192.168.62.26
  set transform-set vpn-test
  match address b1022
!
no spanning-tree vlan 100
!
redundancy
  mode rpr-plus
  main-cpu
    auto-sync running-config
    auto-sync standard
!
interface Loopback0
  description Loopback0
  ip address 10.57.255.251 255.255.255.255
!
interface Tunnel0
  description vpn5-2600-1-0000
  bandwidth 1000000
  ip address 10.60.0.193 255.255.255.252
  ip hold-time eigrp 1 35
  ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
  load-interval 30
  tunnel source 192.168.251.1
  tunnel destination 192.168.0.2
!
interface Tunnel1
  description vpn5-2600-2-0001
  bandwidth 1000000
  ip address 10.60.1.193 255.255.255.252
  ip hold-time eigrp 1 35
  ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
  load-interval 30
  tunnel source 192.168.251.1
  tunnel destination 192.168.1.2
!
interface Tunnel2
```

```

description vpn5-2600-3-0002
bandwidth 1000000
ip address 10.60.2.193 255.255.255.252
ip hold-time eigrp 1 35
ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
load-interval 30
tunnel source 192.168.251.1
tunnel destination 192.168.2.2
!
! . . . repetitive lines omitted . . .
!
interface Tunnel1020
description ci26-2600-11-1020
bandwidth 1000000
ip address 10.67.64.193 255.255.255.252
ip hold-time eigrp 1 35
ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
load-interval 30
tunnel source 192.168.251.1
tunnel destination 192.168.60.26
!
interface Tunnel1021
description ci26-2600-12-1021
bandwidth 1000000
ip address 10.67.65.193 255.255.255.252
ip hold-time eigrp 1 35
ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
load-interval 30
tunnel source 192.168.251.1
tunnel destination 192.168.61.26
!
interface Tunnel1022
description ci26-2600-13-1022
bandwidth 1000000
ip address 10.67.66.193 255.255.255.252
ip hold-time eigrp 1 35
ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
load-interval 30
tunnel source 192.168.251.1
tunnel destination 192.168.62.26
!
interface GigabitEthernet2/1
description GigabitEthernet2/1 Outside Interface
no ip address
load-interval 30
crypto connect vlan 100
!
interface GigabitEthernet4/1
description GigabitEthernet4/1
no ip address
load-interval 30
flowcontrol receive on
flowcontrol send off
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,1002-1005
switchport mode trunk
cdp enable
!
interface GigabitEthernet4/2
description GigabitEthernet4/2
no ip address
load-interval 30
flowcontrol receive on

```



```
flowcontrol send off
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,1002-1005
switchport mode trunk
cdp enable
spanning-tree portfast trunk
!
interface GigabitEthernet5/1
description GigabitEthernet5/1 Inside Interface
ip address 10.57.1.1 255.255.255.0
load-interval 30
!
interface Vlan100
description Vlan100
ip address 192.168.251.1 255.255.255.0
load-interval 30
no mop enabled
crypto map static-map
!
router eigrp 1
network 10.0.0.0
no auto-summary
!
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.251.2
no ip http server
no ip http secure-server
!
ip access-list extended b0000
 permit gre host 192.168.251.1 host 192.168.0.2
ip access-list extended b0001
 permit gre host 192.168.251.1 host 192.168.1.2
ip access-list extended b0002
 permit gre host 192.168.251.1 host 192.168.2.2
ip access-list extended b0003
 permit gre host 192.168.251.1 host 192.168.3.2
!
! . . . reptitive lines omitted . . .
!
ip access-list extended b1020
 permit gre host 192.168.251.1 host 192.168.60.26
ip access-list extended b1021
 permit gre host 192.168.251.1 host 192.168.61.26
ip access-list extended b1022
 permit gre host 192.168.251.1 host 192.168.62.26
!
snmp-server community public RO
snmp-server community private RW
snmp-server system-shutdown
snmp-server enable traps tty
!
alias exec macedon remote command switch test lcp 4 lcp 1
!
line con 0
 exec-timeout 0 0
 password cisco
 login
line vty 0 4
 exec-timeout 0 0
 password cisco
 login
 transport input lat pad mop telnet rlogin udptn nasi ssh
line vty 5 15
```

```

exec-timeout 0 0
password cisco
login
!
ntp clock-period 17179687
ntp server 10.57.1.2
end

```

Cisco 7600/Sup720/VPN SPA Headend Configuration (p2p GRE on Sup720)

In this configuration, the Cisco 7600 platform is aggregating the p2p GRE over IPsec tunnels, with crypto tunnels aggregated to the VPN SPA and p2p GRE being handled by the Sup720.

Headend #1:

```

hostname vpn6-7600-1
!
ip multicast-routing
no ip domain-lookup
ipv6 mfib hardware-switching replication-mode ingress
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
no mls acl tcam share-global
mls cef error action freeze
no scripting tcl init
no scripting tcl enmdir
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
crypto isakmp keepalive 10
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
!
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
!
redundancy
mode sso
main-cpu
  auto-sync running-config
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
power redundancy-mode combined
no diagnostic cns publish
no diagnostic cns subscribe
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
!

```

```
interface Loopback0
  description Loopback0
  ip address 192.168.246.1 255.255.255.255
  load-interval 30
!
interface Loopback1
  description Loopback1
  ip address 192.168.246.2 255.255.255.255
  load-interval 30
!
interface Loopback2
  description Loopback2
  ip address 192.168.246.3 255.255.255.255
  load-interval 30
!
! . . . repetitive lines omitted . . .
!
interface Loopback999
  description Loopback999
  ip address 192.168.249.250 255.255.255.255
  load-interval 30
!
interface Loopback1000
  description Loopback1000
  ip address 10.57.255.251 255.255.255.255
  load-interval 30
!
interface Tunnel0
  description vpn5-2800-1-0000
  bandwidth 1000000
  ip address 10.60.0.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.246.1
  tunnel destination 192.168.0.2
!
interface Tunnel1
  description vpn5-2800-2-0001
  bandwidth 1000000
  ip address 10.60.1.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.246.2
  tunnel destination 192.168.1.2
!
interface Tunnel2
  description vpn5-2800-3-0002
  bandwidth 1000000
  ip address 10.60.2.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.246.3
  tunnel destination 192.168.2.2
!
! . . . repetitive lines omitted . . .
!
interface Tunnel998
  description ci25-2600-19-0998
  bandwidth 1000000
  ip address 10.67.18.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.249.249
  tunnel destination 192.168.38.26
```

```

!
interface Tunnel999
  description ci25-2600-20-0999
  bandwidth 1000000
  ip address 10.67.19.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.249.250
  tunnel destination 192.168.39.26
!
interface GigabitEthernet3/1
  description GigabitEthernet3/1 Outside Interface
  no ip address
  load-interval 30
  crypto connect vlan 100
!
interface GigabitEthernet4/0/1
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 1,100,1002-1005
  switchport mode trunk
  mtu 9216
  no ip address
  flowcontrol receive on
  flowcontrol send off
  spanning-tree portfast trunk
!
interface GigabitEthernet4/0/2
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 1,1002-1005
  switchport mode trunk
  mtu 9216
  no ip address
  flowcontrol receive on
  flowcontrol send off
  spanning-tree portfast trunk
!
interface GigabitEthernet5/1
  description GigabitEthernet5/1 Inside Interface
  ip address 10.57.1.1 255.255.255.0
  no ip redirects
  ip pim sparse-mode
  load-interval 30
!
interface Vlan100
  description Vlan100
  ip address 192.168.241.1 255.255.255.0
  load-interval 30
  no mop enabled
  crypto map dynamic-map
  crypto engine subslot 4/0
!
router ospf 1
  router-id 10.57.255.251
  log-adjacency-changes
  area 10.60.0.0 range 10.60.0.0 255.255.192.0
  area 10.60.64.0 range 10.60.64.0 255.255.192.0
  area 10.60.128.0 range 10.60.128.0 255.255.192.0
  area 10.60.192.0 range 10.60.192.0 255.255.192.0
  area 10.61.0.0 range 10.61.0.0 255.255.192.0
  area 10.61.64.0 range 10.61.64.0 255.255.192.0
  area 10.61.128.0 range 10.61.128.0 255.255.192.0
  area 10.61.192.0 range 10.61.192.0 255.255.192.0

```

```
area 10.62.0.0 range 10.62.0.0 255.255.192.0
area 10.62.64.0 range 10.62.64.0 255.255.192.0
area 10.62.128.0 range 10.62.128.0 255.255.192.0
area 10.62.192.0 range 10.62.192.0 255.255.192.0
area 10.63.0.0 range 10.63.0.0 255.255.0.0
area 10.64.0.0 range 10.64.0.0 255.255.192.0
area 10.64.64.0 range 10.64.64.0 255.255.192.0
area 10.64.128.0 range 10.64.128.0 255.255.192.0
area 10.64.192.0 range 10.64.192.0 255.255.192.0
area 10.65.0.0 range 10.65.0.0 255.255.192.0
area 10.65.64.0 range 10.65.64.0 255.255.192.0
area 10.65.128.0 range 10.65.128.0 255.255.192.0
area 10.65.192.0 range 10.65.192.0 255.255.192.0
area 10.66.0.0 range 10.66.0.0 255.255.192.0
area 10.66.64.0 range 10.66.64.0 255.255.192.0
area 10.66.128.0 range 10.66.128.0 255.255.192.0
area 10.66.192.0 range 10.66.192.0 255.255.192.0
area 10.67.0.0 range 10.67.0.0 255.255.192.0
network 10.57.0.0 0.0.255.255 area 0.0.0.0
network 10.60.0.0 0.0.63.255 area 10.60.0.0
network 10.60.64.0 0.0.63.255 area 10.60.64.0
network 10.60.128.0 0.0.63.255 area 10.60.128.0
network 10.60.192.0 0.0.63.255 area 10.60.192.0
network 10.61.0.0 0.0.63.255 area 10.61.0.0
network 10.61.64.0 0.0.63.255 area 10.61.64.0
network 10.61.128.0 0.0.63.255 area 10.61.128.0
network 10.61.192.0 0.0.63.255 area 10.61.192.0
network 10.62.0.0 0.0.63.255 area 10.62.0.0
network 10.62.64.0 0.0.63.255 area 10.62.64.0
network 10.62.128.0 0.0.63.255 area 10.62.128.0
network 10.62.192.0 0.0.63.255 area 10.62.192.0
network 10.63.0.0 0.0.255.255 area 10.63.0.0
network 10.64.0.0 0.0.63.255 area 10.64.0.0
network 10.64.64.0 0.0.63.255 area 10.64.64.0
network 10.64.128.0 0.0.63.255 area 10.64.128.0
network 10.64.192.0 0.0.63.255 area 10.64.192.0
network 10.65.0.0 0.0.63.255 area 10.65.0.0
network 10.65.64.0 0.0.63.255 area 10.65.64.0
network 10.65.128.0 0.0.63.255 area 10.65.128.0
network 10.65.192.0 0.0.63.255 area 10.65.192.0
network 10.66.0.0 0.0.63.255 area 10.66.0.0
network 10.66.64.0 0.0.63.255 area 10.66.64.0
network 10.66.128.0 0.0.63.255 area 10.66.128.0
network 10.66.192.0 0.0.63.255 area 10.66.192.0
network 10.67.0.0 0.0.63.255 area 10.67.0.0
!
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.241.2
!
no ip http server
ip pim autorp listener
!
snmp-server community public RO
snmp-server community private RW
snmp-server system-shutdown
!
control-plane
!
dial-peer cor custom
!
line con 0
exec-timeout 0 0
password cisco
login
```

```

line vty 0 4
  exec-timeout 0 0
  password cisco
  login
!
ntp server 10.57.1.2
no cns aaa enable
end

```

Cisco 7600/Sup720/VPN SPA Headend Configuration (p2p GRE on VPN SPA)

In this configuration, the Cisco 7600 platform is aggregating the p2p GRE over IPsec tunnels, with both p2p GRE and crypto tunnels aggregated to the VPN SPA.

Headend #1:

```

hostname vpn6-7600-1
!
ip multicast-routing
no ip domain-lookup
ipv6 mfib hardware-switching replication-mode ingress
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
no mls acl tcam share-global
mls cef error action freeze
no scripting tcl init
no scripting tcl encdir
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
crypto isakmp keepalive 10
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
!
crypto map dynamic-map local-address Vlan100
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
!
redundancy
  mode sso
  main-cpu
  auto-sync running-config
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
power redundancy-mode combined
no diagnostic cns publish
no diagnostic cns subscribe
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000

```

```
!
!
interface Loopback0
  description Loopback0
  ip address 10.57.255.251 255.255.255.255
!
interface Tunnel0
  description vpn5-2800-1-0000
  bandwidth 1000000
  ip address 10.60.0.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.0.2
!
interface Tunnel1
  description vpn5-2800-2-0001
  bandwidth 1000000
  ip address 10.60.1.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.1.2
!
interface Tunnel2
  description vpn5-2800-3-0002
  bandwidth 1000000
  ip address 10.60.2.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.2.2
!
! . . . repetitive lines omitted . . .
!
interface Tunnel998
  description ci25-2600-19-0998
  bandwidth 1000000
  ip address 10.67.18.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.38.26
!
interface Tunnel999
  description ci25-2600-20-0999
  bandwidth 1000000
  ip address 10.67.19.193 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.39.26
!
interface GigabitEthernet3/1
  description GigabitEthernet3/1 Outside Interface
  no ip address
  load-interval 30
  crypto connect vlan 100
!
interface GigabitEthernet4/0/1
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 1,100,1002-1005
  switchport mode trunk
```

```

mtu 9216
no ip address
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
interface GigabitEthernet4/0/2
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,1002-1005
switchport mode trunk
mtu 9216
no ip address
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
interface GigabitEthernet5/1
description GigabitEthernet5/1 Inside Interface
ip address 10.57.1.1 255.255.255.0
no ip redirects
ip pim sparse-mode
load-interval 30
!
interface Vlan100
description Vlan100
ip address 192.168.241.1 255.255.255.0
load-interval 30
no mop enabled
crypto map dynamic-map
crypto engine subslot 4/0
!
router ospf 1
router-id 10.57.255.251
log-adjacency-changes
area 10.60.0.0 range 10.60.0.0 255.255.192.0
area 10.60.64.0 range 10.60.64.0 255.255.192.0
area 10.60.128.0 range 10.60.128.0 255.255.192.0
area 10.60.192.0 range 10.60.192.0 255.255.192.0
area 10.61.0.0 range 10.61.0.0 255.255.192.0
area 10.61.64.0 range 10.61.64.0 255.255.192.0
area 10.61.128.0 range 10.61.128.0 255.255.192.0
area 10.61.192.0 range 10.61.192.0 255.255.192.0
area 10.62.0.0 range 10.62.0.0 255.255.192.0
area 10.62.64.0 range 10.62.64.0 255.255.192.0
area 10.62.128.0 range 10.62.128.0 255.255.192.0
area 10.62.192.0 range 10.62.192.0 255.255.192.0
area 10.63.0.0 range 10.63.0.0 255.255.0.0
area 10.64.0.0 range 10.64.0.0 255.255.192.0
area 10.64.64.0 range 10.64.64.0 255.255.192.0
area 10.64.128.0 range 10.64.128.0 255.255.192.0
area 10.64.192.0 range 10.64.192.0 255.255.192.0
area 10.65.0.0 range 10.65.0.0 255.255.192.0
area 10.65.64.0 range 10.65.64.0 255.255.192.0
area 10.65.128.0 range 10.65.128.0 255.255.192.0
area 10.65.192.0 range 10.65.192.0 255.255.192.0
area 10.66.0.0 range 10.66.0.0 255.255.192.0
area 10.66.64.0 range 10.66.64.0 255.255.192.0
area 10.66.128.0 range 10.66.128.0 255.255.192.0
area 10.66.192.0 range 10.66.192.0 255.255.192.0
area 10.67.0.0 range 10.67.0.0 255.255.192.0
network 10.57.0.0 0.0.255.255 area 0.0.0.0
network 10.60.0.0 0.0.63.255 area 10.60.0.0
network 10.60.64.0 0.0.63.255 area 10.60.64.0

```



```
network 10.60.128.0 0.0.63.255 area 10.60.128.0
network 10.60.192.0 0.0.63.255 area 10.60.192.0
network 10.61.0.0 0.0.63.255 area 10.61.0.0
network 10.61.64.0 0.0.63.255 area 10.61.64.0
network 10.61.128.0 0.0.63.255 area 10.61.128.0
network 10.61.192.0 0.0.63.255 area 10.61.192.0
network 10.62.0.0 0.0.63.255 area 10.62.0.0
network 10.62.64.0 0.0.63.255 area 10.62.64.0
network 10.62.128.0 0.0.63.255 area 10.62.128.0
network 10.62.192.0 0.0.63.255 area 10.62.192.0
network 10.63.0.0 0.0.255.255 area 10.63.0.0
network 10.64.0.0 0.0.63.255 area 10.64.0.0
network 10.64.64.0 0.0.63.255 area 10.64.64.0
network 10.64.128.0 0.0.63.255 area 10.64.128.0
network 10.64.192.0 0.0.63.255 area 10.64.192.0
network 10.65.0.0 0.0.63.255 area 10.65.0.0
network 10.65.64.0 0.0.63.255 area 10.65.64.0
network 10.65.128.0 0.0.63.255 area 10.65.128.0
network 10.65.192.0 0.0.63.255 area 10.65.192.0
network 10.66.0.0 0.0.63.255 area 10.66.0.0
network 10.66.64.0 0.0.63.255 area 10.66.64.0
network 10.66.128.0 0.0.63.255 area 10.66.128.0
network 10.66.192.0 0.0.63.255 area 10.66.192.0
network 10.67.0.0 0.0.63.255 area 10.67.0.0
!
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.241.2
!
no ip http server
ip pim autorp listener
!
snmp-server community public RO
snmp-server community private RW
snmp-server system-shutdown
!
control-plane
!
dial-peer cor custom
!
line con 0
  exec-timeout 0 0
  password cisco
  login
line vty 0 4
  exec-timeout 0 0
  password cisco
  login
!
ntp server 10.57.1.2
no cns aaa enable
end
```

Cisco 7200VXR/7600 Dual Tier Headend Architecture Configurations

This configuration is for the Cisco 7200VXR terminating p2p GRE and the Cisco 7600 with Sup720 and VPN SPA providing high-capacity encryption.

```
hostname vpn2-7200-1
```

```

!
boot-start-marker
boot-end-marker
!
logging buffered 65535 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
clock calendar-valid
no aaa new-model
ip subnet-zero
!
ip cef
no ip domain lookup
!
ip ips po max-events 100
no ftp-server write-enable
!
interface Tunnel0
  description vpn5-2800-1-0000
  bandwidth 1000000
  ip address 10.60.0.193 255.255.255.252
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.0.2
!
interface Tunnel1
  description vpn5-2800-2-0001
  bandwidth 1000000
  ip address 10.60.1.193 255.255.255.252
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.1.2
!
interface Tunnel2
  description vpn5-2800-3-0002
  bandwidth 1000000
  ip address 10.60.2.193 255.255.255.252
  load-interval 30
  tunnel source 192.168.241.1
  tunnel destination 192.168.2.2
!
! . . . repetitive lines omitted . . .
!
interface Tunnel998
  description ci25-2600-19-0998
  bandwidth 1000000
  ip address 10.67.18.193 255.255.255.252
  load-interval 30
  tunnel source 192.168.245.1
  tunnel destination 192.168.38.26
!
interface Tunnel999
  description ci25-2600-20-0999
  bandwidth 1000000
  ip address 10.67.19.193 255.255.255.252
  load-interval 30
  tunnel source 192.168.245.1
  tunnel destination 192.168.39.26
!
interface Loopback0
  description Loopback0
  ip address 10.57.255.251 255.255.255.255

```

```
!
interface GigabitEthernet0/1
  description GigabitEthernet0/1
  no ip address
  load-interval 30
  duplex full
  speed 1000
  media-type gbic
  negotiation auto
!
interface GigabitEthernet0/1.241
  description GigabitEthernet0/1.241
  encapsulation dot1Q 241
  ip address 192.168.241.1 255.255.255.0
!
interface GigabitEthernet0/1.242
  description GigabitEthernet0/1.242
  encapsulation dot1Q 242
  ip address 192.168.242.1 255.255.255.0
!
interface GigabitEthernet0/1.243
  description GigabitEthernet0/1.243
  encapsulation dot1Q 243
  ip address 192.168.243.1 255.255.255.0
!
interface GigabitEthernet0/1.244
  description GigabitEthernet0/1.244
  encapsulation dot1Q 244
  ip address 192.168.244.1 255.255.255.0
!
interface GigabitEthernet0/1.245
  description GigabitEthernet0/1.245
  encapsulation dot1Q 245
  ip address 192.168.245.1 255.255.255.0
!
interface GigabitEthernet0/2
  description GigabitEthernet0/2
  ip address 10.57.1.1 255.255.255.0
  load-interval 30
  duplex auto
  speed auto
  media-type gbic
  negotiation auto
!
router ospf 1
  router-id 10.57.255.251
  log-adjacency-changes
  area 0.0.0.0 range 10.56.0.0 255.252.0.0
  area 10.60.0.0 range 10.60.0.0 255.255.192.0
  area 10.60.64.0 range 10.60.64.0 255.255.192.0
  area 10.60.128.0 range 10.60.128.0 255.255.192.0
  area 10.60.192.0 range 10.60.192.0 255.255.192.0
  area 10.61.0.0 range 10.61.0.0 255.255.192.0
  area 10.61.64.0 range 10.61.64.0 255.255.192.0
  area 10.61.128.0 range 10.61.128.0 255.255.192.0
  area 10.61.192.0 range 10.61.192.0 255.255.192.0
  area 10.62.0.0 range 10.62.0.0 255.255.192.0
  area 10.62.64.0 range 10.62.64.0 255.255.192.0
  area 10.62.128.0 range 10.62.128.0 255.255.192.0
  area 10.62.192.0 range 10.62.192.0 255.255.192.0
  area 10.63.0.0 range 10.63.0.0 255.255.0.0
  area 10.64.0.0 range 10.64.0.0 255.255.192.0
  area 10.64.64.0 range 10.64.64.0 255.255.192.0
  area 10.64.128.0 range 10.64.128.0 255.255.192.0
```

```

area 10.64.192.0 range 10.64.192.0 255.255.192.0
area 10.65.0.0 range 10.65.0.0 255.255.192.0
area 10.65.64.0 range 10.65.64.0 255.255.192.0
area 10.65.128.0 range 10.65.128.0 255.255.192.0
area 10.65.192.0 range 10.65.192.0 255.255.192.0
area 10.66.0.0 range 10.66.0.0 255.255.192.0
area 10.66.64.0 range 10.66.64.0 255.255.192.0
area 10.66.128.0 range 10.66.128.0 255.255.192.0
area 10.66.192.0 range 10.66.192.0 255.255.192.0
area 10.67.0.0 range 10.67.0.0 255.255.192.0
network 10.57.0.0 0.0.255.255 area 0.0.0.0
network 10.60.0.0 0.0.63.255 area 10.60.0.0
network 10.60.64.0 0.0.63.255 area 10.60.64.0
network 10.60.128.0 0.0.63.255 area 10.60.128.0
network 10.60.192.0 0.0.63.255 area 10.60.192.0
network 10.61.0.0 0.0.63.255 area 10.61.0.0
network 10.61.64.0 0.0.63.255 area 10.61.64.0
network 10.61.128.0 0.0.63.255 area 10.61.128.0
network 10.61.192.0 0.0.63.255 area 10.61.192.0
network 10.62.0.0 0.0.63.255 area 10.62.0.0
network 10.62.64.0 0.0.63.255 area 10.62.64.0
network 10.62.128.0 0.0.63.255 area 10.62.128.0
network 10.62.192.0 0.0.63.255 area 10.62.192.0
network 10.63.0.0 0.0.255.255 area 10.63.0.0
network 10.64.0.0 0.0.63.255 area 10.64.0.0
network 10.64.64.0 0.0.63.255 area 10.64.64.0
network 10.64.128.0 0.0.63.255 area 10.64.128.0
network 10.64.192.0 0.0.63.255 area 10.64.192.0
network 10.65.0.0 0.0.63.255 area 10.65.0.0
network 10.65.64.0 0.0.63.255 area 10.65.64.0
network 10.65.128.0 0.0.63.255 area 10.65.128.0
network 10.65.192.0 0.0.63.255 area 10.65.192.0
network 10.66.0.0 0.0.63.255 area 10.66.0.0
network 10.66.64.0 0.0.63.255 area 10.66.64.0
network 10.66.128.0 0.0.63.255 area 10.66.128.0
network 10.66.192.0 0.0.63.255 area 10.66.192.0
network 10.67.0.0 0.0.63.255 area 10.67.0.0
!
ip classless
ip route 192.168.0.0 255.255.255.252 192.168.241.2
ip route 192.168.0.4 255.255.255.252 192.168.241.2
ip route 192.168.0.8 255.255.255.252 192.168.242.2
!
! . . . repetitive lines omitted . . .
!
ip route 192.168.159.16 255.255.255.252 192.168.244.2
ip route 192.168.159.20 255.255.255.252 192.168.245.2
!
no ip http server
no ip http secure-server
!
snmp-server community public RO
snmp-server community private RW
snmp-server system-shutdown
snmp-server enable traps tty
!
control-plane
!
gatekeeper
shutdown
!
line con 0
exec-timeout 0 0
password cisco

```

```
login
transport preferred all
transport output all
stopbits 1
line aux 0
transport preferred all
transport output all
stopbits 1
line vty 0 4
exec-timeout 0 0
password cisco
login
transport preferred all
transport input all
transport output all
!
ntp server 10.57.1.2
!
End
```

Cisco 7600/Sup720/VPN SPA Headend Configuration

```
hostname vpn6-7600-1
!
no aaa new-model
clock timezone EST -5
clock summer-time EDT recurring
clock calendar-valid
ip subnet-zero
!
no ip domain-lookup
ipv6 mfib hardware-switching replication-mode ingress
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
no mls acl tcam share-global
mls cef error action freeze
no scripting tcl init
no scripting tcl encdir
!
crypto isakmp policy 10
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
crypto isakmp keepalive 10
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
!
crypto map dynamic-map local-address Vlan100
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
!
redundancy
mode sso
main-cpu
  auto-sync running-config
spanning-tree mode pvst
```

```

no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
power redundancy-mode combined
no diagnostic cns publish
no diagnostic cns subscribe
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
interface GigabitEthernet3/1
description GigabitEthernet3/1 Outside Interface
no ip address
load-interval 30
crypto connect vlan 100
!
interface GigabitEthernet4/0/1
description GigabitEthernet4/0/1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,100,1002-1005
switchport mode trunk
mtu 9216
no ip address
load-interval 30
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
interface GigabitEthernet4/0/2
description GigabitEthernet4/0/2
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,1002-1005
switchport mode trunk
mtu 9216
no ip address
load-interval 30
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
interface GigabitEthernet5/1
description GigabitEthernet5/1 to vpn2-7200-1 GE0/1
ip address 192.168.181.2 255.255.255.0 secondary
ip address 192.168.161.2 255.255.255.0
no ip redirects
load-interval 30
!
interface GigabitEthernet5/2
description GigabitEthernet5/2 to vpn2-7200-2 GE0/1
ip address 192.168.191.2 255.255.255.0 secondary
ip address 192.168.171.2 255.255.255.0
no ip redirects
load-interval 30
!
interface Vlan100
description Vlan100
ip address 192.168.241.1 255.255.255.0
load-interval 30
no mop enabled
crypto map dynamic-map
crypto engine subslot 4/0
!

```

```
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.241.2
!
no ip http server
!
snmp-server community public RO
snmp-server community private RW
snmp-server system-shutdown
!
control-plane
!
dial-peer cor custom
!
line con 0
  exec-timeout 0 0
  password cisco
  login
line vty 0 4
  exec-timeout 0 0
  password cisco
  login
!
ntp server 10.57.1.2
no cns aaa enable
end
```

ISR Branch Configuration

The following shows relevant configurations for one branch router. For resiliency, two tunnels are configured (primary and secondary), one to each headend. The EIGRP delay metric is used to make Tunnel0 the preferred path. This configuration shows QoS for VoIP flows (shaping and queuing) applied to the physical (outside) interface, the recommended use of summary routes, and an EIGRP stub configuration.

Branch #1:

```
hostname vpn5-2800-1-0000
!
boot-start-marker
boot-end-marker
!
logging buffered 32768 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
no network-clock-participate aim 0
no network-clock-participate aim 1
no aaa new-model
ip subnet-zero
!
ip cef
!
ip ips po max-events 100
no ip domain lookup
ip multicast-routing
no ftp-server write-enable
!
class-map match-all VOICE
  match ip dscp ef
class-map match-any CALL-SETUP
```

```

match ip dscp af31
match ip dscp cs3
class-map match-any INTERNETWORK-CONTROL
match ip dscp cs6
match access-group name IKE
class-map match-all TRANSACTIONAL-DATA
match ip dscp af21
!
policy-map 512kb
class CALL-SETUP
bandwidth percent 2
class INTERNETWORK-CONTROL
bandwidth percent 5
class TRANSACTIONAL-DATA
bandwidth percent 22
queue-limit 16
class VOICE
priority 168
class class-default
fair-queue
queue-limit 6
policy-map 512kb-shaper
class class-default
shape average 486400 4864 0
service-policy 512kb
!
crypto isakmp policy 1
encr 3des
authentication pre-share
group 2
crypto isakmp key bigsecret address 192.168.241.1
crypto isakmp keepalive 10
!
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto map static-map local-address Serial0/0/0
crypto map static-map 10 ipsec-isakmp
set peer 192.168.241.1
set transform-set vpn-test
match address b000
!
interface Tunnel0
description Tunnel0
bandwidth 512
ip address 10.60.0.194 255.255.255.252
ip pim sparse-mode
load-interval 30
tunnel source 192.168.0.2
tunnel destination 192.168.241.1
!
interface Loopback0
description Loopback0
ip address 10.60.0.254 255.255.255.255
ip pim sparse-mode
!
interface FastEthernet0/1
description FastEthernet0/1
ip address 10.60.0.129 255.255.255.192 secondary
ip address 10.60.0.1 255.255.255.128
load-interval 30
duplex full
speed 100
!

```



```
interface Serial10/0/0
  description Serial10/0/0
  bandwidth 512
  ip address 192.168.0.2 255.255.255.252
  service-policy output 512kb-shaper
  load-interval 30
  tx-ring-limit 1
  tx-queue-limit 1
  crypto map static-map
!
router ospf 1
  router-id 10.60.0.254
  log-adjacency-changes
  passive-interface FastEthernet0/1
  network 10.0.0.0 0.255.255.255 area 10.60.0.0
!
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.0.1
no ip http server
no ip http secure-server
ip pim autorp listener
!
ip access-list extended IKE
  permit udp any any eq isakmp
ip access-list extended b000
  permit gre host 192.168.0.2 host 192.168.241.1
!
snmp-server community private RW
snmp-server community public RO
snmp-server system-shutdown
snmp-server enable traps tty
!
control-plane
!
line con 0
  exec-timeout 0 0
  password cisco
  logging synchronous
  login
line aux 0
line vty 0 4
  exec-timeout 0 0
  password cisco
  logging synchronous
  login
!
ntp source Loopback0
ntp server 10.57.3.255
!
End
```




Legacy Platform Test Results

This chapter contains scalability and performance data for older Cisco products. The testing has been performed under the same conditions as outlined in the design guide.

Cisco Headend VPN Routers (Legacy)

Table B-1 Cisco Headend Router Platform Throughput with p2p GRE over IPsec (no ACL, PAT, IOS-FW, IPS, or QoS)

Router Platform	Hardware Acceleration	# of Tunnels Active	Throughput Kpps	Throughput Mbps	CPU % Utilization
Cisco 7200VXR with NPE-400	SA-VAM	240	23.5	48.4	80%
Cisco 7200VXR with NPE-300	SA-VAM	240	18.8	38.3	79%
Cisco 3745	AIM-HPH	120	8.7	18.1	78%

Cisco Branch Office VPN Routers (Legacy)

Table B-2 Cisco Branch Router Platform Throughput with p2p GRE over IPsec (1 Tunnel, IOS-FW, PAT, and ACLs but not QoS or IPS)

Branch Router Platform	Hardware Acceleration Option	Throughput Kpps	Throughput Mbps	CPU % Utilization
Cisco 3745	AIM-HPH	13.4	28.7	82%
Cisco 3725	AIM-EPII	6.8	15.5	81%
Cisco 3660	AIM-HPH	4.8	10.9	80%
Cisco 2691	AIM-EPII	5.1	11.4	81%
Cisco 2651XM	AIM-BPII	1.3	3.0	85%
Cisco 1760	VPN-Module	0.9	2.0	81%

Table B-2 Cisco Branch Router Platform Throughput with p2p GRE over IPsec (1 Tunnel, IOS-FW, PAT, and ACLs but not QoS or IPS) (continued)

Cisco 1711	On-board	0.8	2.0	81%
Cisco 831	On-board	0.4	1.0	74%



References and Reading

This section includes the following references for further information:

- Documents (available at the following URL: <http://www.cisco.com/go/srnd>):
 - IPsec VPN WAN Design Overview
 - IPsec Direct Encapsulation Design Guide
 - p2p GRE over IPsec Design Guide
 - Dynamic Multipoint VPN (DMVPN) Design Guide
 - Virtual Tunnel Interface (VTI) Design Guide
 - Voice and Video Enabled IPsec VPN (V3PN) Design Guide
 - IPsec Multicast-Enabled VPN Design Guide
 - HA for IPsec: Redundancy and Load Balancing Design Guide
 - Configuring Digital Certificates Design Guide
- Request For Comment (RFC) papers:
 - Security Architecture for the Internet Protocol—RFC2401
 - IP Authentication Header—RFC2402
 - The Use of HMAC-MD5-96 within ESP and AH—RFC2403
 - The Use of HMAC-SHA-1-96 within ESP and AH—RFC2404
 - The ESP DES-CBC Cipher Algorithm With Explicit IV—RFC2405
 - IP Encapsulating Security Payload (ESP)—RFC2406
 - The Internet IP Security Domain of Interpretation for ISAKMP—RFC2407
 - Internet and Key Management Protocol (ISAKMP)—RFC2408
 - The Internet Key Exchange (IKE)—RFC2409
 - The NULL Encryption Algorithm and Its Use With IPsec—RFC2410
 - IP Security Document Roadmap—RFC2411
 - The OAKLEY Key Determination Protocol—RFC2412



Acronyms

Term	Definition
3DES	Triple Data Encryption Standard
ACL	Access control list
AES	Advanced Encryption Standard
AH	Authentication Header
AIM	Advanced Integration Module
ATM	Asynchronous Transfer Mode
CA	Certificate Authority
CAC	Call Admission Control
CBWFQ	Class-Based Weighted Fair Queuing
CEF	Cisco Express Forwarding
CPE	Customer Premises Equipment
cRTP	Compressed Real-Time Protocol
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DLSw	Data Link Switching
DMVPN	Dynamic Multipoint Virtual Private Network
DMZ	De-Militarized Zone
DNS	Domain Name Service
DPD	Dead Peer Detection
DSL	Digital Subscriber Line
EIGRP	Enhanced Interior Gateway Routing Protocol
ESP	Encapsulating Security Protocol
FIFO	First In First Out
FQDN	Fully Qualified Domain Name
FR	Frame Relay
FRTS	Frame Relay Traffic Shaping
FTP	File Transfer Protocol

GRE	Generic Route Encapsulation
HSRP	Hot Standby Router Protocol
ICMP	Internet Control Message Protocol
IKE	Internet Key Exchange
IOS	Internetwork Operating System
IP	Internet Protocol
IPmc	IP Multicast
IPsec	IP Security
IP GRE	See GRE
ISP	Internet Service Provider
LFI	Link Fragmentation and Interleaving
LLQ	Low Latency Queuing
L2TP	Layer 2 Tunneling Protocol
MDRR	Modified Deficit Round Robin
mGRE	Multipoint Generic Route Encapsulation
MLPPP	Multi-link Point to Point Protocol
MPLS	Multi-Protocol Label Switching
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NetFlow	Cisco IOS component, collects and exports traffic statistics
NHRP	Next Hop Resolution Protocol
NHS	Next-Hop Server
ODR	On-Demand Routing
OSPF	Open Shortest Path First
p2p GRE	Point-to-Point GRE
PAT	Port Address Translation
PBR	Policy-Based Routing
PE	Premises Equipment
PPTP	Point-to-Point Tunneling Protocol
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User System
RTP	Real-Time Protocol
SA	Security Association
SAA	Service Assurance Agent
SHA-1	Secure Hash Algorithm One
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol

SOHO	Small Office/Home Office
SPA	Shared Port Adapter
SRST	Survivable Remote Site Telephony
TCP	Transmission Control Protocol
TED	Tunnel Endpoint Discovery
ToS	Type of Service
UDP	User Datagram Protocol
VAD	Voice Activity Detection
VoIP	Voice over IP
V ³ PN	Voice and Video Enabled IPsec VPN
VAM	VPN Acceleration Module
VPN	Virtual Private Network
VPNSM	VPN Service Module
VPN SPA	VPN Shared Port Adapter
WAN	Wide Area Network
WRED	Weighted Random Early Detection

