



Voice and Video Enabled IPsec VPN (V³PN) Solution Reference Network Design

January 2004

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

Customer Order Number: 956529



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.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study 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, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, 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. (0304R)

Voice and Video Enabled IPsec VPN (V³PN) Solution Reference Network Design
Copyright © 2003, Cisco Systems, Inc.
All rights reserved.



V3PN Solution Reference Network Design Preface ix

About this Publication	ix
Publication Scope	ix
Audience	ix
Obtaining Documentation	x
World Wide Web	x
Documentation CD-ROM	x
Ordering Documentation	x
Documentation Feedback	x
Obtaining Technical Assistance	xi
Cisco.com	xi
Technical Assistance Center	xi
Cisco TAC Web Site	xii
Cisco TAC Escalation Center	xii

CHAPTER 1

V3PN SRND Introduction 1-1

Supporting Designs	1-1
Composite Solution Description	1-2
Solution Benefits	1-3
Solution Scope	1-4
References and Reading	1-4

CHAPTER 2

V3PN Solution Overview and Best Practices 2-1

Solution Overview	2-2
Solution Characteristics	2-4
General Best Practices Guidelines	2-5
General Solution Caveats	2-6

CHAPTER 3

V3PN Solution Components 3-1

IP Telephony (Voice over IP)	3-1
Quality of Service (QoS)	3-2
IP Security (IPSec)	3-4
Issues Specific to V3PN	3-4

Packet Header Overhead Increases 3-5
 cRTP Not Compatible with IPSec 3-5
 Delay Budget 3-5
 Spoke-to-Spoke Crypto Delay 3-5
 FIFO Queue in Crypto Engine 3-6
 Anti-Replay Failures 3-6

CHAPTER 4

Planning and Design 4-1

IP Telephony (Voice over IP) 4-1
 Calculating Delay Budget 4-2
 Hub-to-Spoke versus Spoke-to-Spoke Calling 4-3
 Cisco IP Softphone 4-4
 Quality of Service (QoS) 4-5
 Bandwidth Provisioning for WAN Edge QoS 4-5
 Packet Size—IPSec Encrypted G.729 4-5
 Packet Size—IPSec Encrypted G.711 4-7
 Packet Size—Layer 2 Overhead 4-7
 Special Considerations for Frame Relay Provisioning 4-8
 Bandwidth Allocation by Traffic Category 4-9
 Campus QoS 4-11
 ToS Byte Preservation 4-11
 QoS Pre-Classify 4-12
 IP Security (IPSec) 4-14
 IPSec and GRE Tunnel Design Considerations 4-14
 Firewall Considerations for Transport of VoIP 4-16
 Anti-Replay Considerations 4-16
 Crypto Engine QoS 4-20
 Current VoIP over IPSec Crypto Engine Capabilities 4-20
 LLQ for Crypto Engine 4-21
 When is LLQ for Crypto Engine Required 4-22
 Head-end Topology 4-23
 Head-end Router Locations 4-24
 Service Provider Recommendations 4-24
 Boundary Considerations 4-24
 Cross-Service-Provider Boundaries 4-25
 Service Level Agreements (SLA) 4-26
 Cisco Powered Network References 4-26
 Load Sharing 4-26
 Load Sharing Capabilities 4-27

Encrypted Traffic Appears as a Few, Large Flows	4-27
Minimize Out-of-Order Packets	4-27
Load Sharing Design Approach	4-28
Load Sharing from Head-end to Branch	4-30
Service Provider Considerations for Load Sharing	4-32
E911 and 911 Emergency Services	4-33
Survivable Remote Site Telephony	4-33
Design Checklist	4-35

CHAPTER 5**Product Selection 5-1**

Scalability Test Methodology	5-2
Traffic Profiles	5-3
Additional Voice Quality Validation	5-5
Head-end Product Selection	5-6
Failover and Head-end Availability	5-6
Performance Under Converged V3PN Traffic Profile	5-7
Impact of QoS on VPN Head-end Performance	5-8
Head-End Scalability and Performance Observations	5-9
Branch Office Product Selection	5-9
Product Applicability by Link Speed	5-10
Performance Under Converged V3PN Traffic Profile	5-11
Branch Scalability and Performance Observations	5-14
Network Performance/Convergence	5-15
Software Releases Evaluated	5-17

CHAPTER 6**Implementation and Configuration 6-1**

Routing Protocol, Switching Path and IP GRE Considerations	6-1
Configure Switching Path	6-1
Configure IP GRE Tunnels	6-2
EIGRP Summarization and Network Addressing	6-2
EIGRP hold-time	6-3
IP GRE Tunnel Delay	6-3
QoS Configuration	6-5
Campus QoS—Mapping ToS to CoS	6-5
QoS Trust Boundary	6-6
Configure QoS Class Map	6-6
QoS Policy Map Configuration	6-7
Configuration Example—512 Kbps Branch	6-7

- WAN Implementation Considerations 6-9
 - WAN Aggregation Router Configuration 6-9
 - Frame Relay Traffic Shaping and FRF.12 (LFI) 6-11
 - Attach Service Policy to Frame Relay Map Class 6-14
 - Apply Traffic Shaping to the Output Interface 6-15
 - Applying Service Policy to HDLC Encapsulated T1 Interfaces 6-16
 - Combined WAN and IPSec/IP GRE Router Configuration—Cisco 7200 HDLC/HSSI 6-17
- IKE and IPSec Configuration 6-19
 - Configure ISAKMP Policy and Pre-shared Keys 6-20
 - Configure IPSec Local Address 6-20
 - Configure IPSec Transform-Set 6-21
 - Configure Crypto Map 6-21
 - Apply Crypto Map to Interfaces 6-22
 - Configuring QoS Pre-Classify 6-23
- Implementation and Configuration Checklist 6-24

CHAPTER 7

Verification and Troubleshooting 7-1

- Packet Fragmentation 7-1
- Displaying Anti-Replay Drops 7-2
- Verifying Tunnel Interfaces and EIGRP Neighbors 7-3
- How EIGRP calculates RTO values for Tunnel Interfaces 7-4
- Using NetFlow to Verify Layer-3 Packet Sizes 7-5
- Using NetFlow to Verify ToS Values 7-6
- Sample Show Commands for IPSec 7-8
- Clearing IPSec and IKE Security Associations 7-10
- Sample Show Commands for QoS 7-12

APPENDIX A

Network Diagram Scalability Testbed and Configuration Files A-1

- Head-end VPN Router A-2
- Branch VPN Router—Frame Relay A-5
- Branch VPN Router—HDLC A-8

APPENDIX B

Configuration Supplement—Voice Module, EIGRP Stub, DSCP, HDLC B-1

- Voice Module Configuration B-1
- Router Configuration—vpn18-2600-2 B-3
- Router Configuration—vpn18-2600-3 B-4
- Router Configuration—vpn18-2600-4 B-5

Router Configuration—vpn18-2600-8	B-6
Router Configuration—vpn18-2600-9	B-7
Router Configuration—vpn18-2600-10	B-8
Router Configuration—vpn18-2600-6	B-10

APPENDIX C**Configuration Supplement—Dynamic Crypto Maps, Reverse Route Injection C-1**

INDEX



V³PN Solution Reference Network Design

Preface

This preface presents the following high level sections:

- [About this Publication, page ix](#)
- [Obtaining Documentation, page x](#)
- [Obtaining Technical Assistance, page xi](#)

About this Publication

This section presents two sections:

- [Publication Scope, page ix](#)
- [Audience, page ix](#)

Publication Scope

This Solution Reference Network Design (SRND) publication is intended to provide a set of guidelines for designing, implementing, and deploying Voice and Video Enabled IPsec VPN (V³PN) solutions.

This SRND defines the comprehensive functional components required to build a Site-to-Site Enterprise Virtual Private Network (VPN) solution that can transport IP telephony and video. The Design Guide identifies the individual hardware requirements and their interconnections, software features, management needs, and partner dependencies, to enable a customer deployable, manageable, and maintainable Site-to-Site Enterprise VPN solution.

Audience

This publication is intended to provide guidance to network design specialists, network engineers, telecommunications systems engineers, and data center network managers responsible for integrating Cisco V³PN technology into existing IP infrastructure or building new V³PN-based networking environments.

Content is presented here with the expectation that Cisco Systems Engineers and Customer Support Engineers will use the information provided in combination with internal information to facilitate secure, scalable, and highly available V³PN networks.

Obtaining Documentation

These sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com>

Translated documentation is available at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:
http://www.cisco.com/cgi-bin/order/order_root.pl
- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:
<http://www.cisco.com/go/subscription>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. In the Cisco Documentation home page, click the **Fax** or **Email** option in the “Leave Feedback” section at the bottom of the page. You can e-mail your comments to bug-doc@cisco.com. You can submit your comments by mail by using the response card behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

If you want to obtain customized information and service, you can self-register on Cisco.com. To access Cisco.com, go to this URL:

<http://www.cisco.com>

Technical Assistance Center

The Cisco Technical Assistance Center (TAC) is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two levels of support are available: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Cisco TAC inquiries are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

The Cisco TAC resource that you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

You can use the Cisco TAC Web Site to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to this URL:

<http://www.cisco.com/tac>

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

<http://www.cisco.com/register/>

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC Web Site, you can open a case online by using the TAC Case Open tool at this URL:

<http://www.cisco.com/tac/caseopen>

If you have Internet access, we recommend that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.



V³PN SRND Introduction

This publication extends the Cisco Architecture for Voice, Video, and Integrated Data (AVVID) by enabling voice and video applications to be transported over a site-to-site IPsec VPN. Just as enterprise implementers expect to run these applications over a private WAN, such as Frame Relay or ATM, they also expect to run voice and video across their VPN implementation with the same quality and level of service. Further, the enterprise implementer should be able to do so and have the VPN be fairly transparent to these applications.

To provide these capabilities, Cisco designed Voice and Video Enabled IPsec VPN (V³PN), which integrates three core Cisco technologies: IP Telephony, Quality of Service (QoS), and IP Security (IPsec) VPN. The result is an end-to-end VPN service that can guarantee the timely delivery of latency-sensitive applications such as voice and video.

This chapter presents the following topics:

- [Supporting Designs, page 1-1](#)
- [Composite Solution Description, page 1-2](#)
- [Solution Benefits, page 1-3](#)
- [Solution Scope, page 1-4](#)
- [References and Reading, page 1-4](#)

Supporting Designs

V³PN is designed to overlay non-disruptively on other core Cisco AVVID designs, including:

- *Enterprise Site-to-Site IPsec VPN Design*
Guidelines—http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns142/networking_solutions_design_guidances_list.html
- *Enterprise IP Telephony Design*
Guidelines—http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidances_list.html
- *Enterprise QoS Design*
Guidelines—http://www.cisco.com/application/pdf/en/us/guest/netsol/ns17/c649/ccmigration_09186a00800d67ed.pdf

This SRND will not cover each of these three technologies in detail, but will instead focus on the intersection of, integration of, and interactions between these functions of the network. Familiarity with design and implementation guides for these underlying technologies will be extremely beneficial to the reader. Please review these guides before attempting to implement a V³PN.

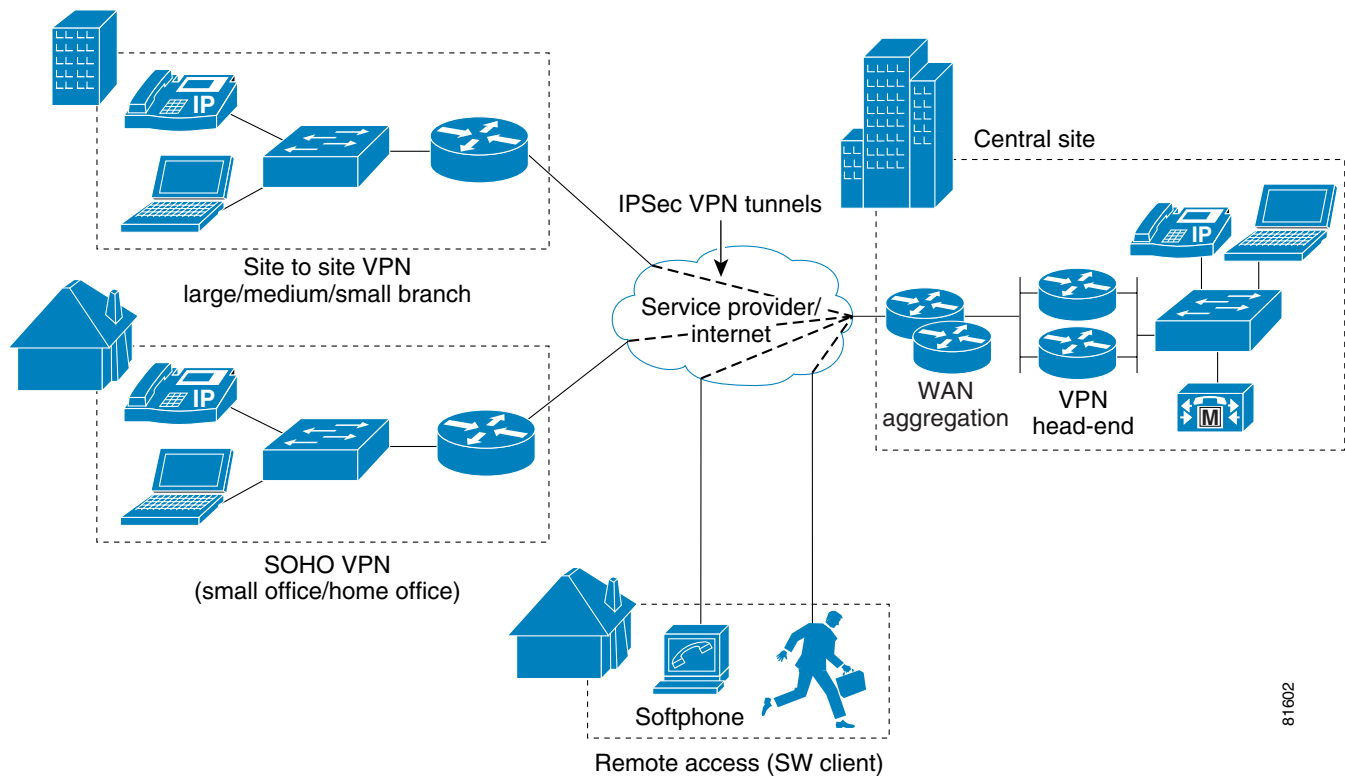
The underlying VPN design principles are based on the SAFE VPN Architecture, therefore the reader should also first be familiar with that architecture and recommendations. Cisco SAFE documentation can be found at: <http://www.cisco.com/go/safe>.

Technical Assistance Center (TAC) Technical Tips are a valuable source of configuration examples for the technologies deployed in this design guide. Please refer to the Technical Tip section after logging on the Cisco TAC Cisco.Com page at: <http://www.cisco.com/tac>.

Composite Solution Description

IPSec VPNs have been deployed as private WAN alternatives for enterprise networks whether managed by the enterprise themselves or as part of a service provider managed service. Figure 1-1 illustrates the composite IPSec VPN deployment models that are deployed today:

Figure 1-1 Composite IPSec VPN Deployment Models



81602

Site-to-site IPSec VPN's are used to connect small, medium, and large branch offices to a central location or locations. This model is referred to in Cisco Enterprise Solutions Engineering Design Guides as *Site-to-Site Branch VPN*.

IPSec VPN's can also be used to connect small office/home office (SOHO) locations to corporate locations. When the VPN connections are *static* (fixed) in nature this model is referred to as *site-to-site SOHO VPN*.

Finally, when the VPN connections are *dynamic* (session-by-session) this model is referred to as *Remote Access VPN*.

The site-to-site branch VPN model is capable of enabling voice and video transport across the VPN in a high quality manner—including transport over service provider networks that support QoS. The site-to-site SOHO VPN model is also capable of transporting high quality voice and video over the VPN; however, broadband service providers are in the early stages of providing QoS support.

This version of this Design Guide focuses primarily on the site-to-site branch VPN model of deployments, as this model currently has the highest level of proven deployability both in terms of Cisco IOS VPN Router functionality as well as service providers being capable of offering a multi-service VPN service to enterprise implementers.

The primary objectives for this Design Guide will be to:

- Define the safe boundaries in which this solution may be deployed including design and implementation considerations as well as highlighting appropriate caveats.
- Provide hardware platform and software code recommendations based on the requirements of a given deployment, including performance and configuration information where applicable.

Since an IPSec VPN deployment involves a service provider this document will delineate requirements of the enterprise as well as what the service provider must provide in order to ensure a successful V³PN deployment.

Solution Benefits

V³PN provides the following benefits for enterprise networks:

- **Higher Productivity**—Enables extension of central site voice, video, and data resources and applications at all corporate sites, thereby enabling employees to work as productively and efficiently as if they were located at the central site.
- **Ease of Provisioning**—V³PN provides enterprises with a flexible means of deploying additional sites that are voice enabled by simply connecting to a service provider instead of procuring private WAN connectivity.
- **Lower Cost**—Pricing for connection via a local Internet service provider is distance insensitive, analogous to Frame Relay. Further, an enterprise can attain converged inter-site connectivity, lowering both the costs of bandwidth and toll cost.
- **Flexibility**—V³PN provides support for extensions to the enterprise applications, such as IP Call Centers (IPCC), Video Conferencing, e-Learning, and Teleworking, irrespective of the physical location of resources and users of these resources.
- **Increased Security**—V³PN is implemented using IPSec encryption and device authentication, thereby providing a higher level of security compared to typical unencrypted and unauthenticated time-division multiplexing (TDM) and voice/video transport.
- **Return on Investment**—Because V³PN is implemented across the Cisco IOS VPN Router product portfolio, existing investments are preserved and can be extended.

For service providers, V³PN provides the following benefits:

- **New Revenue** – Enabling voice and video transport across IPSec VPN's provide a potential source of incremental revenue for the service provider if deploying a Managed Service. The service provider also benefits even if the enterprise manages the IPSec VPN carrying VoIP where as the service provider can achieve incremental revenue by providing value add QoS enabled services.
- **Service Differentiation**—V³PN provides the ability to encrypt voice and video, which is a new security feature that can be offered relative to traditional TDM networks.

- New Customers—By qualifying for the *IP Multi-service* VPN Cisco Powered Network designation, service providers are better positioned to receive new enterprise customers being referred by Cisco account teams for V³PN services.
- Customer Retention—By adding additional value to the enterprise customer, the retention likelihood is greater for service providers, particularly as the enterprise customer becomes more reliant upon the V³PN service for mission critical applications beyond data transport, in other words voice and video.

Solution Scope

This publication will be extended and updated over time as capabilities expand the addressable market. This version of this Design Guide focuses on the following:

- Site-to-site IPsec branch VPN deployment model, where the interface to the service provider is typically a media such as Point-to-Point (PPP), High-Level Data Link Control (HDLC), Frame Relay (FR), Asynchronous Transfer Mode (ATM) or Ethernet (in the case of Metropolitan Area Networks). This Design Guide will be extended in a future revision to include information on the site-to-site SOHO VPN deployment model, typically utilizing DSL or Cable media.
- Cisco IOS VPN Routers to terminate the IPsec VPN tunnels. The PIX platforms will be addressed in a later Design Guide.
- Video and IP Multicast are not fully addressed in this design guide version however where appropriate known design recommendations will be made for both. Tested design recommendations for Video and IP Multicast will be the focus of a subsequent revision of this design guide.
- V³PN was evaluated in a design utilizing IPsec with GRE to support dynamic routing protocols and IP Multicast. However, the performance and scalability results for IPsec/GRE should also be applicable to an IPsec only configuration. An IPsec-only configuration is used as the design for an internal Cisco deployment of V³PN.

Other features that were not evaluated for this revision of the Design Guide include:

- IPsec Stateful Failover
- LZS Compression
- GRE Tunnel Keepalives
- Voice Activity Detection (VAD)

References and Reading

Table 1-1 IETF Requests for Comment

IETF Request for Comment (RFC)	Topic
RFC2401	Security Architecture for the Internet Protocol
RFC2402	IP Authentication Header
RFC2403	The Use of HMAC-MD5-96 within ESP and AH
RFC2404	The Use of HMAC-SHA-1-96 within ESP and AH
RFC2405	The ESP DES-CBC Cipher Algorithm With Explicit IV

Table 1-1 IETF Requests for Comment

IETF Request for Comment (RFC)	Topic
RFC2406	IP Encapsulating Security Payload (ESP)
RFC2407	The Internet IP Security Domain of Interpretation for ISAKMP
RFC2408	Internet and Key Management Protocol (ISAKMP)
RFC2409	The Internet Key Exchange (IKE)
RFC2410	The NULL Encryption Algorithm and Its Use With IPsec
RFC2411	IP Security Document Roadmap
RFC2412	The OAKLEY Key Determination Protocol

Table 1-2 Reference Websites

Topic	Link
Enterprise VPNs	http://www.cisco.com/go/evpn
Cisco SAFE Blueprint	http://www.cisco.com/go/safe
Cisco Network Security	http://www.cisco.com/go/security
Cisco VPN Product Documentation	http://www.cisco.com/univercd/cc/td/doc/product/vpn/
Download VPN Software from CCO	http://www.cisco.com/kobayashi/sw-center/sw-vpn.shtml
Improving Security on Cisco Routers	http://www.cisco.com/warp/public/707/21.html
Essential Cisco IOS Features Every ISP Should Consider	http://www.cisco.com/warp/public/707/EssentialIOSfeatures_pdf.zip
Cisco Technical—Security	http://www.cisco.com/cgi-bin/Support/browse/index.pl?i=Technologies&f=1408
IPSec Support Page	http://www.cisco.com/cgi-bin/Support/PSP/psp_view.pl?p=Internetworking:IPSec
Networking Professionals Connection	http://forums.cisco.com
Voice and Video Enabled IPSec VPN (V ³ PN) Overview	http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns241/netbr09186a00800b0da5.html
Voice and Video Enabled IPSec VPN (V ³ PN) Solution	http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns241/networking_solutions_package.html
NetFlow	http://www.cisco.com/go/netflow



V³PN Solution Overview and Best Practices

This chapter presents a high-level overview of V³PN to give the reader a quick reference as to the capabilities of this solution. The remainder of this document will then go into an increasing level of detail on planning, design, product selection, and implementation of a V³PN.

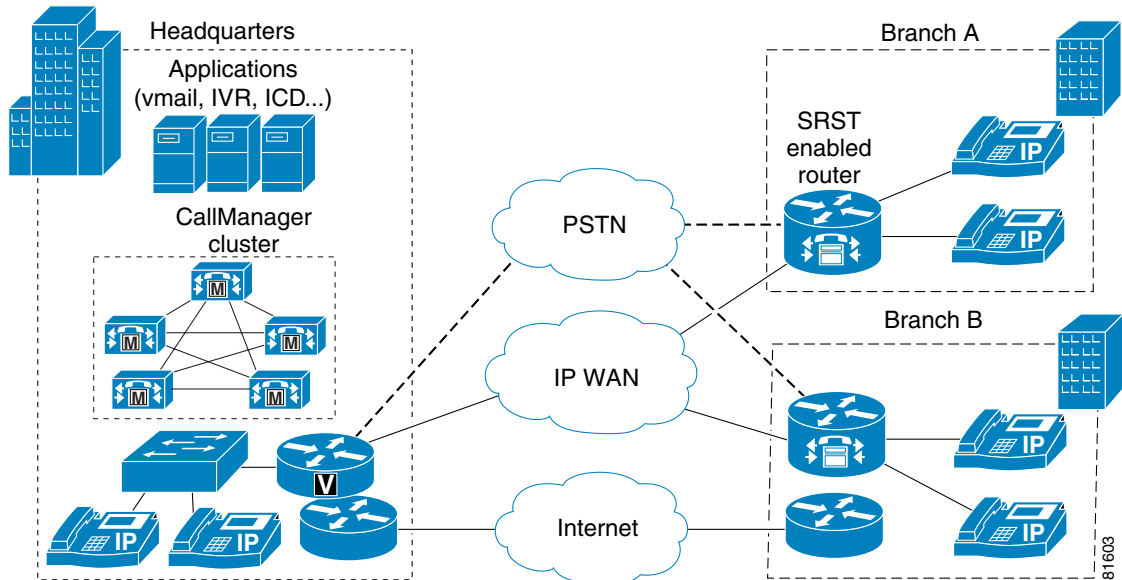
Specific topics in this chapter are:

- [Solution Overview, page 2-2](#)
- [Solution Characteristics, page 2-4](#)
- [General Best Practices Guidelines, page 2-5](#)
- [General Solution Caveats, page 2-6](#)

Solution Overview

Figure 2-1 depicts a typical deployment of IP Telephony using a Centralized Call Processing model.

Figure 2-1 IP Telephony Over Private WAN

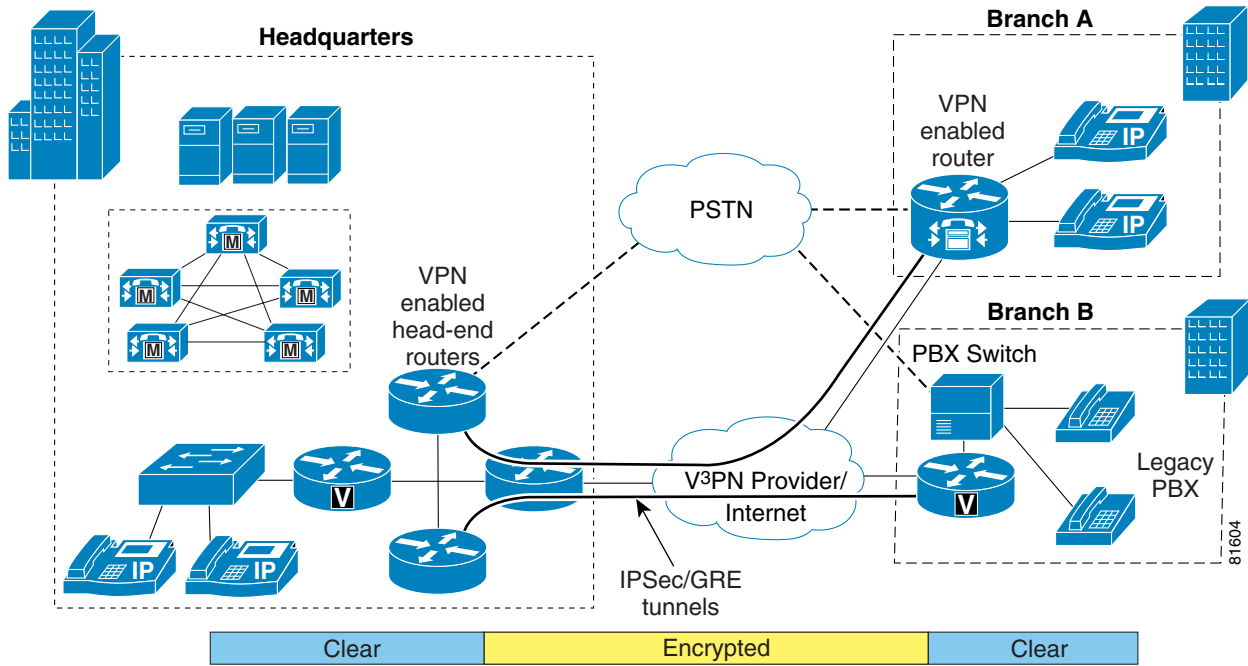


In this arrangement, a Call Manager cluster is deployed at the large central location. Branches might deploy Survivable Remote Site Telephony (SRST), which provides for local call processing to the PSTN in the case of loss of connection to the central site. The PSTN links can also be used for local off-net calling. Connectivity between the central site and branch locations is over a Private WAN technology like FR or ATM. Signaling traffic (such as H.323) is sent over the Private WAN links to the Call Manager cluster at the central site. Voice conversations are established and bearer traffic also flows over the Private WAN links.

Typically, the large central site and large branch locations will have separate connections to an ISP to provide Internet access to the corporation.

Figure 2-2 illustrates this same enterprise implementation with its IP Telephony deployment using a V³PN strategy.

Figure 2-2 IP Telephony Over V³PN



Notice that the IP telephony deployment remains unchanged. But in this deployment, connectivity between the central site and branch locations is over a V³PN provider. Signaling traffic (such as H.323) is sent *encrypted* over the VPN (IPSec/GRE) tunnels to the Call Manager cluster at the central site. Voice conversations are established and bearer traffic also flows *encrypted* over the VPN tunnels.

The encryption provided by IPSec provides an additional level of security for voice conversations. However, neither the IP Phones, Call Manager Cluster, or voice applications such as voice mail servers are aware, nor need to be aware, that their traffic is being transported over a VPN tunnel and being encrypted during transport. The VPN is transparent to these applications.

Another advantage is that typically the V³PN service provider can offer a Layer-3 IP *pipe* such that both the VPN services and Internet services can exist over the single connection to each location. This reduces recurring connection costs as well as reduces the number of devices required for deployment.

Solution Characteristics

Table 2-1 presents the general solution characteristics for V³PN deployments.

Table 2-1 V³PN Solution Characteristics Summary

Solution Characteristics

Secure Triple Data Encryption Standard (3DES) voice, video, and data traffic can be simultaneously transported over the same IPsec VPN tunnels with QoS enabled for high priority traffic, similar to a private WAN, such as Frame Relay and ATM.

Based on Cisco IOS VPN Routers for resiliency, high availability, and a building block approach to high scalability that can support thousands of branch offices.

Scalability and performance evaluation was performed with IPSEC and GRE tunnels although the performance numbers in this document can also be used as a conservative guideline for IPsec only deployments.

The VPN tunnels can be managed by the enterprise or offered by the service provider as a managed service.

IP Telephony traffic traversing an IPsec VPN is transparent to all users and personnel managing the IP Telephony network.

Standard IP Telephony features, such as SRST and different Codec types, are preserved and still possible over V³PN.

Admission control for IP Telephony is handled the same for VPN tunnels as would be for a Private Frame Relay PVC connecting two branch offices together where admission control is based on the max VoIP traffic permitted across a given IPsec tunnel.

Integrated branch routers providing service provider/Internet connection, VPN tunnel termination, IP Telephony gateway, and Cisco IOS Firewall functionality are possible.

General Best Practices Guidelines

Table 2-2 presents a list of *best practices* that have been established through a combination of design experience, scalability and performance evaluation, and internal Cisco trials.

Table 2-2 V³PN Solution Best Practices Guidelines Summary

Solution Best Practices

Deploy hardware-accelerated Encryption on all platforms that support them. SW-based encryption adds unacceptable latency and jitter that significantly degrade voice quality.

Hub-and-spoke IPsec topology is recommended (partial meshing is also possible).

Maximum of 240 (120 active, 120 backup) IP GRE tunnels per head-end router were evaluated, due to the size of the current testbed. Future tests will evaluate up to 480 branches.

Target maximum CPU utilization on each router not to exceed percent that under test maintained EIGRP adjacency on all IP GRE tunnels during failure testing.

IPsec with GRE tunnels are required if IP Multicast or routing protocols (using IP Multicast) is required.

QoS Pre-Classify must be enabled on VPN devices where applicable to ensure appropriate QoS criteria of the encrypted packet can be applied on the egress WAN interface. See the “[QoS Pre-Classify](#)” section on page 4-12 for more information.

G.729 (20 msec sampling at 50 pps) is recommended due to bandwidth consumption after IPsec and GRE overhead are added to the voice packets. See the “[Bandwidth Provisioning for WAN Edge QoS](#)” section on page 4-5 for more information.

Select appropriate Cisco IOS VPN Router products per scalability and performance requirements, as well as link speed that will be deployed. See [Chapter 5, “Product Selection”](#) for more information.

Branch VPN Routers will typically provide both QoS and VPN tunnel termination on an integrated device, while Head-end VPN Routers will typically have a device providing service provider link termination and QoS that is separate from the VPN tunnel termination device.

Use a Cisco Powered Network service provider designated as an *IP Multi-service VPN* provider to ensure the high priority voice and video traffic can be prioritized across the service provider’s network. See the “[Service Provider Recommendations](#)” section on page 4-24 for more information.

Enterprises that have tunnels that traverse multiple service providers must ensure that QoS markings (ToS, IP Precedence or DSCP) and prioritization are preserved and honored when crossing SP boundaries.

Seek a service level agreement (SLA) from the service provider that meets the enterprise organization’s needs in terms of end-to-end delay, jitter and dropped packets. This is analogous to an SLA an enterprise would arrange with a private WAN provider offering Frame Relay or ATM service.

General Solution Caveats

Table 2-3 presents a list of caveats for the solution.

Table 2-3 V³PN Solution Implementation Caveats

Solution Caveats

Compressed RTP (cRTP) and IPSec are incompatible standards. The RTP header is already encrypted when the packet reaches the cRTP engine and therefore cannot be compressed. Industry standardization bodies are currently considering alternative bandwidth optimization techniques for encrypted tunnels.

In Cisco IOS software releases prior to 12.2(13)T, the IPSec Crypto Engine has a FIFO entrance queue. In Cisco IOS software release 12.2(13)T and higher, an LLQ for the IPSec Crypto Engine is supported. See the “[Crypto Engine QoS](#)” section on page 4-20 for more information.

The majority of voice traffic was simulated RTP streams using the NetIQ Chariot test tool, however a real CallManager and IP phones were configured and used for verification.

SRST was implicitly verified, but no performance and scalability evaluation was performed. SRST is supported in Cisco 12.2(7)T on the Cisco 2600 and Cisco 3600 platforms, and on the Cisco 175x series routers in Cisco IOS 12.2(4)XW.

QoS Pre-Classify is supported in 12.2(4)YB on the 1700 series. This is a Business Unit (BU) *special* Cisco IOS software release. All other branch platforms evaluated were running T-train Cisco IOS software.

Multilink PPP was found to have limitations, causing some platforms to drop into process switching and thereby reducing performance. These limitations are being addressed in a future Cisco IOS software release. Frame Relay and HDLC were verified, others such as ATM were not.

At the central site, QoS can be configured on either the VPN head-end devices or on a separate service provider/Internet link terminating device. While both configurations are supported, it is recommended to have separate devices at the central site for scalability.

On branch products (Cisco 2600, Cisco 3600, and Cisco 3700 VPN routers), voice cards and VPN hardware-acceleration cards (AIM) are not supported in 12.2(8)T when installed in the same router. This capability is supported in 12.2(11)T.

The Cisco 806 VPN Router does not support SRST, QoS Pre-Classify, or hardware-accelerated encryption.

The Cisco 806 VPN Router had performance limitations with latency and jitter. It is not a recommended platform for V³PN.



V³PN Solution Components

Implementation of a site-to-site IPSec VPN design capable of supporting transport of voice and video, requires the combination of three Cisco technologies:

- [IP Telephony \(Voice over IP\), page 3-1](#)
- [Quality of Service \(QoS\), page 3-2](#)
- [IP Security \(IPSec\), page 3-4](#)

These three technologies have been implemented on many enterprise networks as standalone functions or in some combination—especially IP Telephony and QoS. V³PN combines all three technologies enabled simultaneously on a common network.

This design guide addresses the areas of intersection between the three technologies and provides configuration and verification tips to promote the successful implementation of a similar design—while maintaining the same availability and voice quality demonstrated in Cisco Enterprise Solutions Engineering lab testing.

IP Telephony (Voice over IP)

This design document assumes the target enterprise site includes or will include an IP Telephony deployment—with the expectation of extending this deployment to branch office locations over an IPSec VPN. As with private WAN deployments of IP Telephony, Call Admission Control (CAC) is implemented to limit the number of concurrent calls based on the capabilities and traffic handling capacity of the branch office routers being deployed, as well as the link speed being used.

IP Telephony can be deployed using several different CODEC and sampling rate schemes. Each offers advantages and disadvantages in terms of voice quality, added latency, bandwidth consumption, and router resource consumption. For example, G.711 with 20 msec sampling and a transmission rate of 50 pps is a common deployment. This offers high voice quality, minimal latency, but higher bandwidth consumption. G.729 with 20 msec sampling and a transmission rate of 50 pps is another common deployment, especially for lower speed links, as it offers very good voice quality at a lower bandwidth consumption.

Both G.711 and G.729 voice calls were evaluated as part of this solution. G.729 with 20 msec sampling and a transmission rate of 50 pps is the recommended CODEC/sampling scheme for V³PN deployments. This is primarily due to the bandwidth requirements.

Another possibility would be to deploy G.729 with 30 msec sampling at a transmission rate of 33 pps. This offers an additional bandwidth savings, although there might be a trade-off in voice quality as loss of a single packet can produce an audible *click* or *pop*. With 20 msec sampling, loss of two consecutive packets would be required to cause audible errors.

Voice Activity Detection (VAD) is another IP Telephony feature that can be used to achieve bandwidth savings. VAD works by not transmitting the natural periods of silence during a voice call, lowering the transmission rate (in packets per second) during these periods and hence lowering the overall bandwidth consumption (on average). Use of VAD was not evaluated.

**Note**

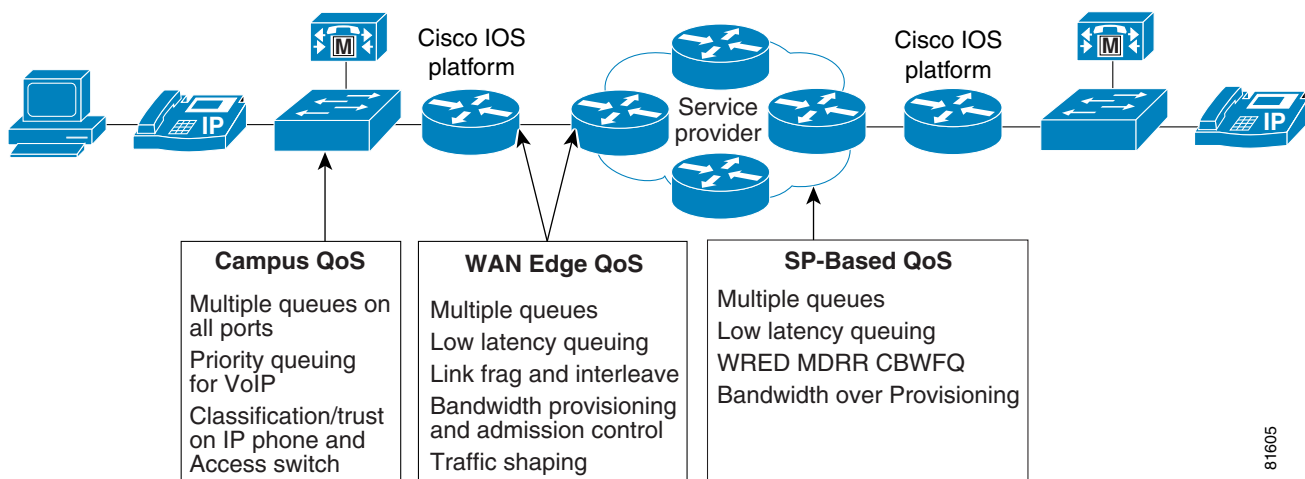
Use of VAD has the potential to result in perceived voice quality issues and is not recommended. For more information, refer to *Enterprise IP Telephony Design Guidelines* at the following location: http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidance_s_list.html

This Design Guide does not cover IP Telephony deployment in any detail, only the components that affect V³PN deployment. Refer to *Enterprise IP Telephony Design Guidelines*.

Quality of Service (QoS)

Cisco QoS is an enabling technology for IP Telephony that prioritizes latency-sensitive packets, such as voice and video, over lower priority traffic (i.e. data) to minimize end-to-end delay. QoS also seeks to insure consistent voice packet delivery which minimizes the arrival variance (jitter). In order for VoIP over IPsec to be successfully implemented, QoS is required at different points in the network. Figure 3-1 shows the different QoS components that are required:

Figure 3-1 Components of QoS for V³PN Deployment



81605

The QoS requirements for the central site (campus) LAN and branch LAN are the same as a typical IP Telephony deployment, including:

- Classification or Class of Service (CoS)/Trust on IP Phone and switch
- Proper speed and duplex between IP Phone and switch
- Multiple queues on IP Phone and switch ports
- WRED within data queue for congestion management
- Mapping ToS to CoS on the LAN edge router

On the enterprise WAN edge, QoS requirements are again the same as for a typical IP Telephony deployment, including:

- CBWFQ with a priority queue (LLQ)
- Traffic Shaping (if applicable)
- Link Fragmentation and Interleaving (LFI) (where appropriate)

One major difference is that now the bandwidth provisioning at the WAN edge must consider the additional overhead of IPSec and IP GRE. Another major difference could be if Compressed Real-Time Protocol (cRTP) is currently being used. IPSec and cRTP are not compatible standards, therefore cRTP is not applicable for V³PN implementations. Encrypted IPSec packets are ignored by cRTP logic, therefore no bandwidth savings will result for the encrypted voice calls.

This Design Guide does not cover campus or WAN edge QoS deployment in any detail, only the components that affect V³PN deployment. Refer to *the Enterprise QoS Design Guidelines* at:

http://www.cisco.com/application/pdf/en/us/guest/netsol/ns17/c649/ccmigration_09186a00800d67ed.pdf

The majority of scalability and performance evaluation was based on Frame Relay and HDLC as the service provider's WAN technology. The natural progression of this solution is to enable V³PN across the Internet. QoS enabled ISP's are in the early adoption phase of deployment. Some ISPs offer SLAs for end-to-end latency and packet drop rates. However, these agreements might treat all data traffic equally, with no distinction between packets with different IP Precedence or differentiated service code point (DSCP) values within the ToS byte.

The requirements for the service provider include:

- Deliver packets end-to-end with minimal delay, jitter, and loss. This can be accomplished by prioritizing packets based on ToS (IP Precedence/DSCP) in the core (QoS enabled core) or over-provisioning the bandwidth in the core
- Provide and meet a SLA to achieve end-to-end delay, jitter, and packet loss requirements. Several SPs offer SLAs with less than 60 msec delay today.
- Implement a policy to handle high priority traffic exceeding the agreed upon rate with the enterprise organization, as well as how traffic will be handled which crosses the service provider boundary.
- Highly preferred that the service provider *mirror* the enterprise WAN edge QoS: CBWFQ, LLQ, Traffic Shaping and LFI.

The service provider QoS component is an important aspect of voice call quality—one over which the enterprise network implementer has the least direct control. For this reason, the ability of a service provider to offer Cisco Powered Network *IP Multi-service VPN* is a competitive advantage and point of service differentiation in the competitive ISP market.

A service provider offering different classes of service within their backbone would need to police the organization's data rate by ToS byte to the agreed upon offered rate, or charge more for transport of the higher priority traffic. Without this policing or tiered billing function, their customer could (intentionally or inadvertently) mark all packets going into the service provider with the highest priority.

For more information regarding service providers for V³PN and SLA requirements, see the “[Service Provider Recommendations](#)” section on page 4-24.

IP Security (IPSec)

The IPSec component provides secrecy (confidentially) and integrity of both voice and data over public networks. Government regulations might legislate the use of crypto in financial and health care enterprises, but the driving motivation is ubiquitous low cost network access by Internet Service Providers.

This publication does not cover site-to-site IPSec VPN deployment in absolute detail, only the components that affect V³PN deployment. Refer to the *Enterprise Site-to-Site IPSec VPN Design Guide* for more information at the following location:

http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns142/networking_solutions_design_guidances_list.html

The design principles were proven by scalability testing in the Cisco Enterprise Solutions Engineering lab. The scale testing methods are designed to test worst-case scenarios. From a design standpoint that entails:

- Strong (3DES) encryption for both IKE and IPSec
- IP GRE with IPSec Tunnel mode
- Routing protocol (EIGRP)
- Diffie-Hellman Group 2 (1024 bit) for IKE
- Secure Hash Algorithm (SHA)-HMAC, a 160-bit rather than 128-bit with Message Digest 5 (MD5)-HMAC (both hash algorithms are truncated to 12 bytes in the ESP packet trailer).

Should the organization choose to implement less stringent security parameters, or choose to use IPSec Transport mode rather than Tunnel mode, or not implement IP GRE tunnels, the test results will continue to be applicable from a scalability standpoint.

Advanced Encryption Standard (AES) is a new encryption algorithm, available now in several Cisco platforms. AES will be evaluated in a future revision.

Pre-shared keys were used during this evaluation. Digital Certificates are planned for subsequent scalability, however they are currently deployed in internal Cisco deployments. Whether pre-shared keys or Digital Certificates are being used are not expected to significantly affect V³PN performance.

Issues Specific to V³PN

Combining voice, video, and data traffic on a converged network, encrypting that traffic with IPSec, then prioritizing that traffic with QoS presents unique design considerations for the network designer. The sections that follow address these issues in more detail.

- [Packet Header Overhead Increases](#)
- [cRTP Not Compatible with IPSec](#)
- [Delay Budget](#)
- [Spoke-to-Spoke Crypto Delay](#)
- [FIFO Queue in Crypto Engine](#)
- [Anti-Replay Failures](#)

Packet Header Overhead Increases

The addition of an IP GRE header and IPsec / ESP header increases the size of the original voice (or video) packet. Using Layer 3 packet sizes, a 60-byte G.729 voice packet increases to 136 bytes with IP GRE and IPsec tunnel mode. A 200-byte G.711 voice packet increases to 280 bytes. For low-speed WAN links, G.729 has been the recommended CODEC to conserve bandwidth. This is considerably more important in an IPsec implementation.

See the [“Bandwidth Provisioning for WAN Edge QoS”](#) section on page 4-5 for more information regarding IPsec and GRE packet expansion.

cRTP Not Compatible with IPsec

Network managers implement cRTP as a link efficiency mechanism to decrease the overhead of voice traffic on low-speed (less than E1) links. cRTP compresses the IP/UDP/RTP header in an RTP data packet from 40 bytes to approximately 2-to-5 bytes. However, cRTP and IPsec are inherently incompatible standards. The original IP/UDP/RTP header is encrypted by IPsec by the time the RTP compressor is called upon to perform the compression. Therefore cRTP cannot associate the encrypted IP/UDP/RTP packet with a known media stream, compression does not occur, and bandwidth savings are not realized. The encrypted IP/UDP/RTP packet simply by-passes the compression process and continues.

Delay Budget

The delay budget for a typical IP Telephony implementation includes fixed and variable components. The ITU target value is for one-way delay to be 150 msec or less, although up to 250 msec might be acceptable.

In a V³PN deployment, there are two additional delay components added in to the overall delay budget: Encryption at the origination point of the VPN tunnel and decryption at the termination point. Performance and scalability testing results suggest that in most cases the additional delay caused by encryption/decryption is approximately 2-to-5 msec. A conservative planning estimate would be 10 msec for encryption and 10 msec for decryption. Fundamentally, IPsec adds a trivial amount of additional delay to voice deployments versus a clear text IP telephony deployment.

See the [“Calculating Delay Budget”](#) section on page 4-2 for more information regarding delay budget.

Spoke-to-Spoke Crypto Delay

In addition to the aforementioned delay components, and the crypto engine encrypt and decrypt delay component, this design guide specifies a hub-and-spoke topology. If one branch places a voice call to another branch, the voice packet must be encrypted at the branch, decrypted at the head-end router, the routing table lookup (path determination) switches the packet out another IP GRE tunnel, which is encrypted and then decrypted at the second branch. The frequency of branch-to-branch calling must be analyzed and if the call frequency warrants, requires creation of a site-to-site VPN tunnel between these particular branch locations (partial meshing).

Dynamic Multipoint VPN (DMVPN) provides an alternative to hub-and-spoke topologies—essentially providing a virtual fully-meshed VPN. This feature can be of great benefit in V³PN designs by reducing the dependence of VoIP traffic flowing from a spoke (branch) through the central site to another spoke (branch). DMVPN allows a dynamic tunnel to be established directly between the two branch locations. More information regarding DMVPN will be added to this document in a future revision.

See the [“Calculating Delay Budget” section on page 4-2](#) for more information regarding delay budget.

FIFO Queue in Crypto Engine

In Cisco IOS software releases prior to 12.2(13)T, the crypto engine is a single FIFO queue for admitting packets to the crypto process. Both packets for encryption and packets for decryption share the same queue. In some hardware platforms and traffic profiles, the crypto engine could be the gating throughput factor. LLQ for Crypto Engine is available in Cisco IOS software release 12.2(13)T and higher, to manage any over subscription of the hardware crypto accelerator. On platforms and traffic profiles tested in this guide, voice quality was maintained without LLQ for Crypto Engine.

Software-based crypto engines (with no HW-accelerators installed) have a FIFO queue and will not have the LLQ for Crypto Engine feature. For this reason, as well as the unacceptable latency and jitter performance, software-based crypto is not recommended for V³PN implementations.

See the [“Crypto Engine QoS” section on page 4-20](#) for more information regarding Crypto Engine and QoS.

Anti-Replay Failures

The IPSec ESP (Encapsulating Security Protocol) authentication component provides message integrity. One aspect is a sequence number assigned to packets within each security association. This sequencing of packets prevents a packet from being captured and replayed at a later time to the intended receiver. Fundamentally QoS techniques alter the order of packets between two IPSec peers, voice packets are prioritized over data packets. The message integrity aspect of IPSec runs contrary to the intended re-ordering of packets by QoS. Under the scenarios evaluated for this guide, there was no significant packet loss due to anti-replay/QoS interaction.

See the [“Anti-Replay Considerations” section on page 4-16](#) or more information regarding IPSec anti-replay and its interaction with QoS.



Planning and Design

This chapter addresses planning and design considerations for enabling V³PN. It reviews issues and design considerations specific to IP Telephony, QoS and IPSec. Specifics on product selection for branch and head-end devices are also provided for review. An overview on service provider considerations is also provided. The following specific planning and design sections are presented in this chapter:

- [IP Telephony \(Voice over IP\), page 4-1](#)
- [Quality of Service \(QoS\), page 4-5](#)
- [IP Security \(IPSec\), page 4-14](#)
- [Head-end Topology, page 4-23](#)
- [Head-end Router Locations, page 4-24](#)
- [Service Provider Recommendations, page 4-24](#)
- [Load Sharing, page 4-26](#)
- [E911 and 911 Emergency Services, page 4-33](#)
- [Survivable Remote Site Telephony, page 4-33](#)

This chapter ends with the “[Design Checklist](#)” section on [page 4-35](#) to further facilitate V³PN planning.

IP Telephony (Voice over IP)

In this solution, IP Telephony can be thought of as an application transported on a site-to-site IPSec VPN. As such, it is an application with special requirements – more stringent than most data applications – and thus bears special consideration. These requirements include:

- Packets arrive at a constant rate (assuming no VAD)¹
- Packets arrive in *per call* increments (do not have a portion of a call)
- Quality of the call is a function of jitter, latency and packet loss
- Call Admission Control must be addressed—same as with a Frame Relay deployment

For planning purposes the packet arrival rate is assumed to be 50 packets per second (pps), per call. A call is assumed to be 50 pps transmitted and received. During solution testing, jitter, latency and packet loss are monitored and reported in test results as a gauge of expected voice quality. Like IP Telephony

1. The design throughout this document assumes the Voice Activity Detection (VAD) feature of IP Telephony is disabled. VAD has far-reaching implications on a design and resulting voice quality that are beyond the scope of this solution.

deployments over a private WAN, there is still a requirement to provide Call Admission Control (CAC) as additional voice calls (over-subscription) cannot be permitted to impact the voice quality of established calls.

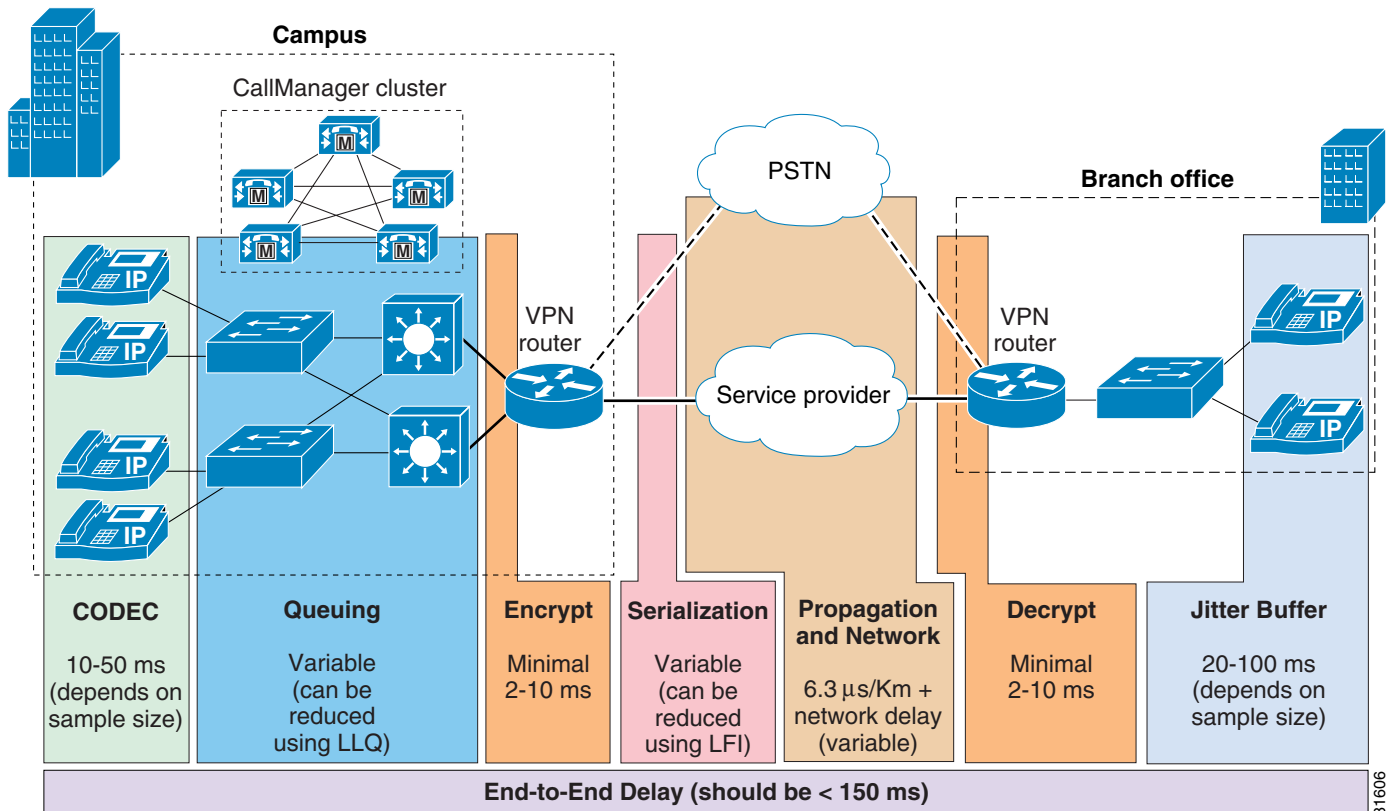
Calculating Delay Budget

One important network design aspect of implementing IP Telephony involves the calculation of the end-to-end, one-way delay budget. Ideally the total budget will fall under 100 msec. At 100 msec, there is the possibility that the phone will be answered but the called party will not hear the caller's greeting. This is referred to as *signal delay*, *voice clipping*, or *voice path cut-through delay*.

The ITU standard G.114 states that a one-way delay budget of 150 msec is acceptable for high voice quality. For most networks, a delay of 200 msec will provide acceptable voice quality with 250 msec as the upper limit.

An additional delay component when running IP Telephony over IPsec VPN is the delay of encrypting and decrypting the voice packets. This is especially important when voice calls are from branch to branch in a hub and spoke environment, as all but the coder and dejitter components are duplicated for each spoke traversed. The specific end-to-end components of the delay budget are shown in Figure 4-1.

Figure 4-1 Calculating End-to-End, One-Way Delay Budget



In most deployments, the addition of a few milliseconds of additional delay for encryption and decryption is insignificant when compared to the total delay budget. For all platforms that support hardware crypto accelerators, they are highly recommended for voice deployments. Software encryption can introduce unacceptable latency and jitter as the CPU becomes fully utilized, which significantly

degrades voice quality. For example, issuing a *copy running start (write memory)* can cause a CPU spike for a few seconds and degrade voice quality. Hardware crypto accelerators help minimize intermittent voice quality issues associated with software crypto and CPU spikes.

**Note**

For planning purposes, 2-to-5 msec of additional delay is added for encryption and decryption, under normal conditions (no over-subscription of the crypto engine). Refer also to the “[Crypto Engine QoS](#)” section on page 4-20.

One additional consideration is the codec/sampling scheme being deployed. This delay component can generally be quantified by considering the sampling duration plus any compression delay. For example, G.711 with 20 msec sampling has a delay of approximately 20 msec—only the sampling delay plus less than one msec of encoding and processing delay. In comparison, G.729 with 20 msec sampling has a delay of approximately 25 msec, the sampling delay of 20 msec plus approximately 5 msec for compression, encoding, and processing.

In the Cisco Enterprise Solutions Engineering lab testing, Chariot endpoints report the end-to-end delay of the voice (RTP) streams, as well as jitter and packet loss. Chariot reported values do not include encoding, packetization and dejitter buffer delay. For testing purposes, target threshold values reported by Chariot were as follows:

- Voice jitter—Under 10 msec
- Voice delay—Under 50 msec
- Voice loss—Under 0.5 percent loss

See chapter “[Scalability Test Methodology](#)” section on page 5-2 for more information regarding the test methodology and thresholds established for product performance determination.

**Note**

It is also important for Call Signaling packets to experience minimal delay across the network, or call setup issues can result. Refer to *Enterprise IP Telephony Design Guidelines* for more information (http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidances_list.html).

Hub-to-Spoke versus Spoke-to-Spoke Calling

Hub and Spoke topologies are prevalent in most enterprise networks. A typical design includes one or two corporate data centers and perhaps a *hot* business recovery site. The remote locations are tied into both data centers and either a third consolidation point connects to the business recovery site or the remote offices have a direct third connection to the business recovery site. Traffic flow, sources and sinks of data, are to and from the data centers and individual remote locations. There is very little traffic volume between the remote offices, little spoke-to-spoke communication. This model is true of hospitality/hotel enterprises, banking, retail or any service oriented business that has large numbers of branch locations which operate in a fairly standalone fashion.

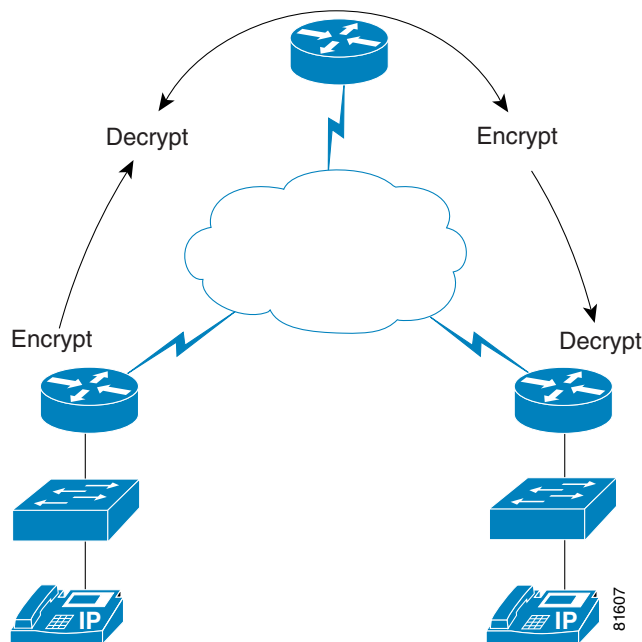
The Cisco Enterprise Solutions Engineering lab testing mimics this design, as a major focus is to identify the scalability of the head-end routers in terms of number of branch locations at various WAN rates.

However, the addition of VoIP as an overlay application can change the hub to spoke paradigm. Consider an implementation where branch locations are now individual knowledge worker’s home offices and whole workgroups or teams are working remotely. VoIP traffic might be spoke-to-spoke, while the data traffic might continue to be hub-to-spoke.

Figure 4-2 illustrates this situation. In this case, spoke-to-spoke IP Telephony must be considered in regards to the delay budget. The *coder* and *dejit* values of the delay budget do not change, but the intermediate delay components discussed in the “Calculating Delay Budget” section on page 4-2 are now duplicated: spoke to head-end, and head-end to spoke. In this example, the VoIP packets are encrypted at the first branch, decrypted at the head-end, routed to another tunnel interface, encrypted again and then decrypted at the receiving spoke. All of the queuing, serialization and network delay is also present on both *legs* of the call.

Both Cisco Enterprise Solutions Engineering lab testing and internal Cisco deployments have demonstrated encrypted spoke-to-spoke voice calls are both feasible and practical provided the overall delay budget is within tolerances.

Figure 4-2 Spoke to Spoke Calling



Dynamic Multipoint VPN (DMVPN) provides an alternative to hub-and-spoke topologies, essentially providing a virtual fully-meshed VPN. This feature can be of great benefit in V³PN designs by allowing a dynamic tunnel to be established directly between the two branch locations. More information regarding DMVPN will be added to this document in a future revision.

Cisco IP Softphone

Cisco IP Softphone is a Windows-based application for the PC. During solution testing a Cisco IP Softphone was included in the test bed and calls were placed between the Softphone and a 7960 phone, from branch to campus. For the purposes of this solution, there are no special considerations which must be addressed if Softphones are deployed in addition to, or in place of, 7960 phones. The Cisco IP Softphone provides the same IP Precedence/DSCP markings as a 7960 for the voice bearer and call setup streams.

**Note**

It should be noted that the Softphone application running on a laptop might be a *best effort* implementation, because the laptop operating system has no provision for QoS. Therefore, it is possible that other applications on the laptop can interfere with voice quality.

Quality of Service (QoS)

IP Telephony deployments have been the catalyst for enhancements in the development and deployment of QoS services in today's networks. Implementing V³PN introduces two new aspects to the traditional QoS implementation to support VoIP:

- Encryption increases the bandwidth requirements of both voice and data, impacting the provisioning of service policies on output interfaces
- Encryption provides confidentiality of portions of the original IP packet that previously were referenced by the output QoS service policy.

These two key QoS issues—and a review of important QoS concepts implemented in this design guide—are addressed in the subsequent sections.

Bandwidth Provisioning for WAN Edge QoS

This section details the bandwidth requirements of encrypted voice, how to provision the enterprise WAN edge to ensure encrypted voice quality.

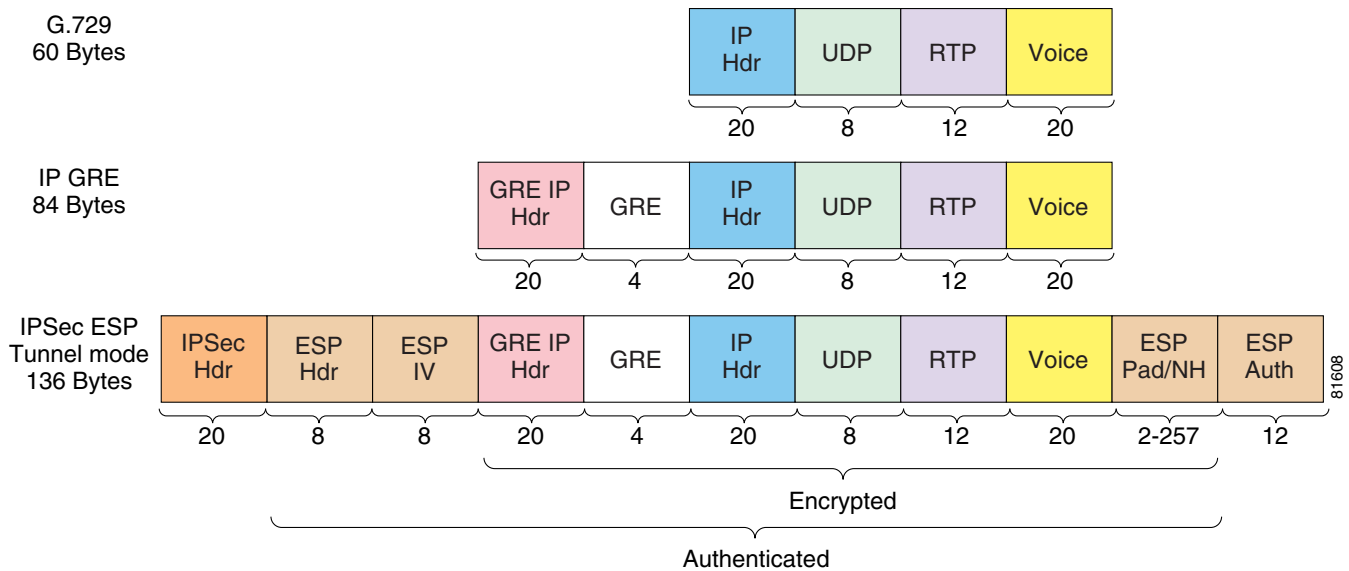
Packet Size—IPSec Encrypted G.729

The Layer 3 data rate for a G.729 call (50 pps) is 24 Kbps. Encrypting that packet using IPSec Tunnel mode for IP GRE increases that rate to approximately 56 Kbps (in each direction). The calculation is as follows:

- **136 bytes** per packet at 50 packets per second = 6,800 bytes or **54,400 bits per second**

The 136-byte packet's header, data and trailer fields are shown below in [Figure 4-3](#).

Figure 4-3 IPsec Encrypted G.729 Packet Anatomy



Start with a 60-byte G.729 IP/UDP/RTP voice packet. IP GRE adds 24 bytes, including a new IP header and GRE encapsulation.

The ESP header contains a 4-byte Security Parameter Index (SPI) field and the 4 byte sequence number. This sequence number is used by ESP authentication—the anti-replay logic.

ESP might add up to 255 bytes of padding. DES is a block cipher, encrypting blocks of 8 bytes (64-bits) at a time, and thus the need for the encryption algorithm to add padding to the plain text. The ESP Authentication Data field must align on a 4-byte boundary. The ESP Pad length field is 1 byte and starts at the third byte of a 4-byte word, and the ESP Next Header field occupies the fourth byte. The Next Header field is used to identify the payload's protocol. The ESP Authentication Data field contains either the MD5 16 byte hash or the SHA-1 20 byte hash, both truncated to 12 bytes. See RFC 2104 regarding truncation of the hash value.

The ESP IV (Initialization Vector) ensures the uniqueness cipher text if the same plain text characters are encrypted in different blocks or messages. It is used by block chaining ciphers like DES. The ESP IV byte count can be determined from the command **show crypto ipsec sa** for the security association (SA). The ESP IV field is considered part of the payload and might not be shown by a protocol analyzer. This is also true of the trailer fields: ESP Padding; Pad Length; Next Header; and Authorization.

An increase of the original packet by one byte might result in no increase in the resulting encrypted packet, or it might increase more than one byte. NetFlow can be used to verify the Layer 3 length of the encrypted and decrypted packet sizes. See the [“Using NetFlow to Verify Layer-3 Packet Sizes” section on page 7-5](#) for an example.

In this example, IPsec adds 52 bytes to the IP GRE packet, so the resulting packet combined with IP GRE and IPsec is 136 bytes.

The G.729 codec family is specified by these four specific designations: g729r8, g729ar8, g729br8, and g729abr8. All generate the same format code word, but differ in complexity and support of voice activity detection (VAD):

- g729r8—This is the default codec, it and all forms of G.729 generate 8,000 bps
- g729ar8—Codec is a simplified version of g729r8
- g729br8—Same as g729r8 but includes VAD

- g729abr8—Simplified g729r8 but includes VAD

With VAD, silence is not sent over the network—only speech, and therefore the total bandwidth used by voice traffic might be much less than the 50 pps specified. Studies have shown that 35 percent bandwidth savings can be realized by the use of VAD. Voice quality might be degraded slightly and comfort noise should be considered to provide audible feedback to the listener that the other party is still on the call.

However, for bandwidth capacity planning purposes and for testing, assume VAD is disabled and the RTP streams are 50 packets per second continuously.

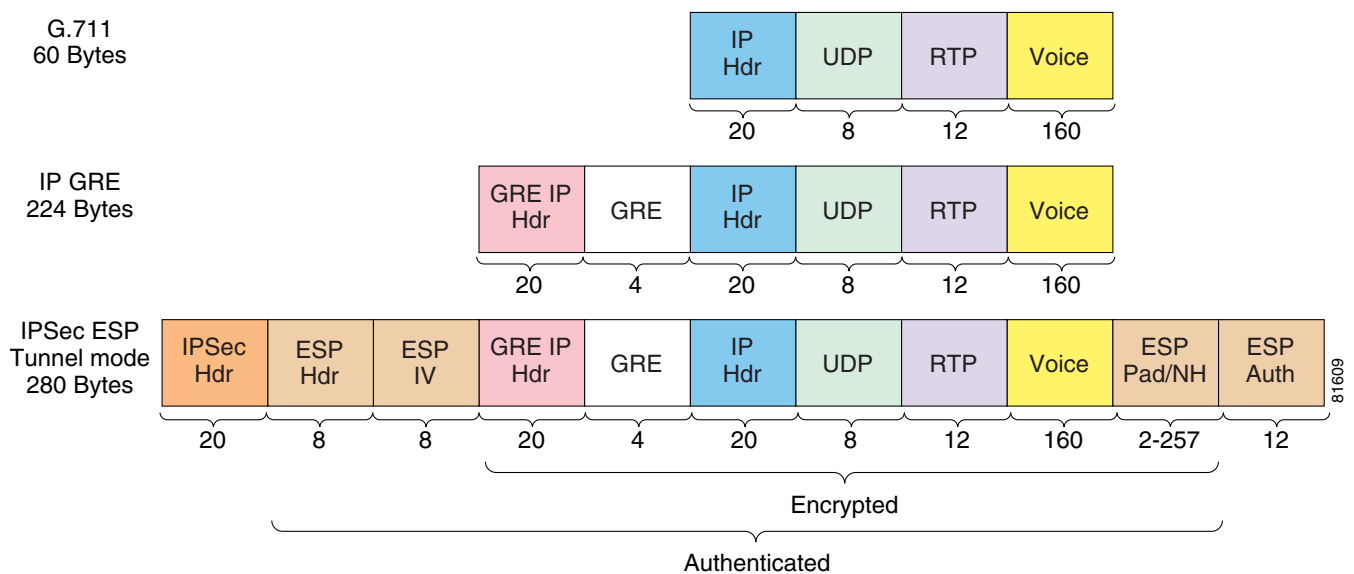
Packet Size—IPSec Encrypted G.711

The Layer 3 data rate for a G.711 call (50 pps) is 80 Kbps. Encrypting that packet using IPSec Tunnel mode for IP GRE increases that rate to approximately 112 Kbps (in each direction). The calculation is as follows:

- **280 bytes** per packet at 50 packets per second = 14,000 bytes or **112,000 bits per second**

The 280 byte packet's header, data and trailer fields are shown in [Figure 4-4](#).

Figure 4-4 *IPSec Encrypted G.711 Packet Anatomy*



The construction of the G.711 encrypted voice packet is the same as the G.729 example in the “[Packet Size—IPSec Encrypted G.729](#)” section on page 4-5. The major difference is that G.711 starts with a 200-byte voice packet. After the addition of IP GRE and IPSec, the packet size becomes 280 bytes.

The codec choices for G711 are g711alaw and g711ulaw and both generate data at 64,000 bits per second, the same data rate as clear-channel. There is no VAD option available with G711.

Packet Size—Layer 2 Overhead

While Layer 2 headers do not increase the original voice packet size to the extent of adding IP GRE and IPSec headers and trailers, they must be included in the calculation for priority (LLQ) queue in the voice class of the CBWFQ policy map. Some common Layer 2 encapsulations are shown in [Table 4-1](#).

Table 4-1 Summary of Common Layer 2 Encapsulations

L2 Encapsulation	Bytes to include bandwidth calculation
Ethernet	14 bytes
Frame Relay	4 bytes ¹
PPP	4 bytes
Multilink PPP	10 bytes
HDLC	4 bytes
ATM	53-byte cell, 48 bytes of payload

1. Frame Relay flags do not need explicit provisioning for the LLQ. For non-fragmented (voice) frames 4 bytes per frame Layer-2 overhead; for fragmented (non-LLQ packets over the frame-relay fragment size) frames, 6 bytes.

To continue the calculation of the G.729 voice packet with IP GRE and IPSec as shown previously, add 4 bytes for the Frame Relay header:

- 136 bytes + 4 bytes = 140 bytes * 50 pps = 7,000 bytes or **56,000** bps

For the G.711 example:

- 280 bytes + 4 bytes = 284 bytes * 50 pps = 14,200 bytes or **113,600** bps

These should be considered the minimum values for bandwidth planning.

Jitter in the path of the voice packets can increase or decrease the arrival rate—for short periods of time, the bit per second values can be slightly higher than calculated above. In the organization-specific environment, the service policy should be reviewed for drops in the voice class using this command:

```
show policy-map interface serial 0/0.100
```

In the event drops are encountered, increase the **priority** keyword value in the voice class to eliminate voice drops. When the interface is not congested, the priority (voice) class can exceed its bandwidth and not be dropped. When the interface is congested, the offered rate of voice traffic is dropped if it exceeds its allocation.

Special Considerations for Frame Relay Provisioning

The primary configuration used for the scalability and performance evaluation was Frame Relay encapsulation, although HDLC was also evaluated.

To minimize serialization delay on low speed links, Link Fragmentation and Interleaving (LFI) must be implemented by Layer 2. For Frame Relay, this is accomplished by Frame Relay Forum's FRF.12 Implementation Agreement (also known as FRF.11 Annex C).

To reserve bandwidth for, and prioritize voice packets, Class Based Weighted Fair Queuing (CBWFQ) is configured with a priority or low-latency queue (LLQ). The CBWFQ service policy is applied to the Frame Relay Traffic Shaping (FRTS) map-class. FRTS is a prerequisite for FRF.12, but it also provides congestion notification to CBWFQ. CBWFQ allocated bandwidth among the configured classes of traffic during periods of interface congestion.

In this guide, the Cisco IOS Frame Relay Traffic Shaping target CIR values is 95 percent of the carrier's configured CIR. While following this guideline adds an extra step to the planning and configuration process, the conservative approach is to implement this best practice. To make configuration easier, tables have been provided for the link speeds under test.

Bandwidth Allocation by Traffic Category

In order to implement a QoS Service Policy it is necessary to decide how traffic will be sub-divided into various categories and assigned relative priorities for those categories.

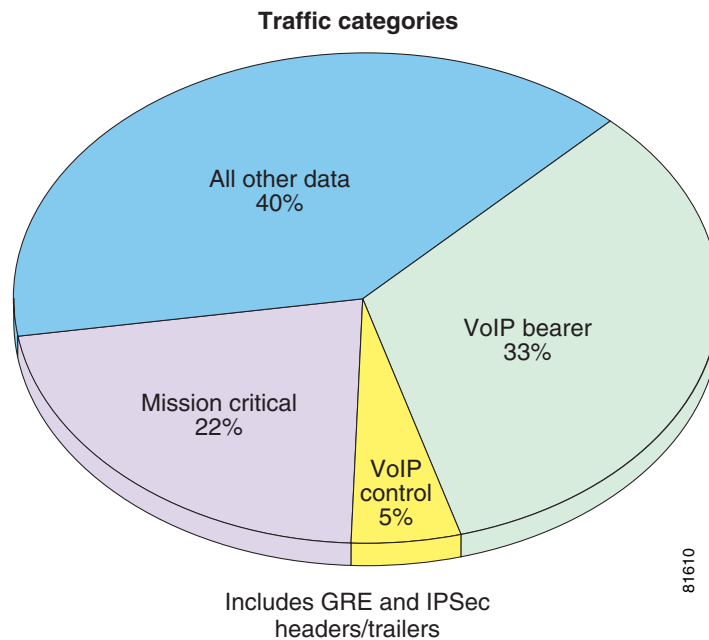
For networks already using QoS, established Service Policies must be considered and might be customized by the enterprise as well.

This design assumes voice and data traffic falls into the categories:

- VoIP control (signaling)
- VoIP bearer (RTP streams)
- Mission critical data traffic
- Other data traffic and overhead

Considerations must be given to the amount of high priority voice traffic allocated on a converged network. This design establishes an upper bound of 33 percent for such traffic, as this places a significant small-packet burden on the VPN. The relative target bandwidth percentages for traffic categories is shown in [Figure 4-5](#).

Figure 4-5 Traffic Categorization for Bandwidth Provisioning



Note

These percentages (and categories) are not strict design rules, they are provided as a guideline. When the bandwidth of a particular class is not being used, it is available to other traffic categories. These percentages simply call out reserved bandwidth percentages.

These allocations must include:

- Payload (such as voice)
- IP GRE overhead
- IPSec overhead

- Layer 2 encapsulation header overhead

With the addition of IP GRE headers and IPSec headers and trailer, and the Layer 2 header, the per call value using a G.729 codec is 56 Kbps.

Using a 512Kbps link as an example, these traffic categories would be configured as:

```
!
policy-map llq-branch
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority 168
  class class-default
    fair-queue
!
```

Voice bandwidth must be configured in per call increments (such as 56 Kbps). Therefore, for lower speed links, the allocation exceeds the 33 percent target.

Class-default queues first-in-first-out (FIFO) unless weighted fair queue (WFQ) is configured. Weighted random early detection (WRED) could also optionally be configured for the mission-critical and class-default. WRED is IP Precedence aware and serves to intelligently drop packets and provide feedback to TCP based applications to reduce their sending data rate. Additionally, WRED configured in a class provides a degree of visibility to the network manager as to the IP Precedence of the traffic that is tail or random dropped.

The key QoS considerations for this configuration are:

- When configuring *percent* in a policy-map, the *percent* value is a percent of the underlying (link or Frame Relay Traffic Shaping MINCIR) bandwidth.
- For serial HDLC-encapsulated interfaces, the sum of the Kbps specified by the classes in the policy map cannot exceed 75 percent of the available bandwidth. This can be manually changed by the **max-reserved-bandwidth** interface command.
- The priority (LLQ) is policed if there is congestion and it exceeds its bandwidth allocation.
- Bandwidth that is not allocated to a class is available to the default class, *class-default*.
- Packets not selected for a class are placed in class-default, regardless of their IP Precedence; class-default is not assumed to be only for packets with IP Precedence of '0'.

It is convenient when configuring a CBWFQ policy-map to specify bandwidth in terms of percentages rather than actual values (in Kbps), since the same **policy-map** can be applied to a range of link speeds without modifications. The sample configuration template for this design guide shows the LLQ specified in Kbps and the *call-setup* and *mission-critical* classes in percent. Using Kbps for the LLQ is easier to illustrate than percentages as bandwidth allocation is in multiples of 56 Kbps per call.

Table 4-2 lists the key QoS and Call Admission Control parameters in this design.

Table 4-2 Bandwidth Allocation Parameters by Link Speed

Line Rate (Kbps)	Maximum Number of G.729 Calls	Max G.729 Calls as Percentage of Line rate	Priority 56k per Call (Kbps)	Call Setup 5 percent (Kbps)	Mission Critical 22 percent (Kbps)	Max-reserved-bandwidth
64	1	87.5	56	3	None	100
128	1	43.7	56	6	26	75
256	2	43.7	112	12	53	75

Table 4-2 Bandwidth Allocation Parameters by Link Speed

Line Rate (Kbps)	Maximum Number of G.729 Calls	Max G.729 Calls as Percentage of Line rate	Priority 56k per Call (Kbps)	Call Setup 5 percent (Kbps)	Mission Critical 22 percent (Kbps)	Max-reserved-bandwidth
512	3	33	168	24	106	75
768	4	29	224	36	160	75
1024	6	33	336	48	213	75
1536	9	33	504	72	320	75
2048	12	33	672	102	450	75

These values should be substituted into the configuration sample to create the branch router configuration, based on the line rate that services the remote location. For branches that connect at 64 Kbps, 128 Kbps and 256 Kbps, the voice traffic exceeds the target of approximately 33 percent voice traffic on the link. These are highlighted, as is the absence of a *mission critical class* and override of the *max-reserved-bandwidth* for 64 Kbps links.

When planning the bandwidth required for a branch office, consider the number of concurrent calls traversing the WAN this branch is expected to make during peak call periods. This varies based on the job function of the employees located at a branch. For example, an office of software engineers would be expected to make fewer calls than an office of telemarketers. One rule of thumb is one call for every six people (1:6), but this could range from 1:4 to 1:10. Given the 512 Kbps link as an example, with a target of 3 G.729 Calls, that link could theoretically support between 12 and 30 people.

If the maximum number of calls is not active on the link, the bandwidth is available for data traffic. If there is a misconfiguration in the call admission control for the branch and more calls are attempted (the voice class exceeds its calculated data rate), CBWFQ will police (drop) packets during congestion and all calls will exhibit poor voice quality. Either CallManager “Locations” or Gatekeeper Call Admission Control, must be implemented to guarantee voice quality.

More information on configuration of CallManager can be found in *Enterprise IP Telephony Design Guidelines* at:

- http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidances_list.html

Campus QoS

Refer to the recommendations on Campus Switching Designs for Cisco AVVID. Campus QoS techniques have been well documented and are not included in this design guide. *Campus QoS Design Guidelines* can be found at:

- http://www.cisco.com/application/pdf/en/us/guest/netsol/ns17/c649/ccmigration_09186a00800d67ed.pdf

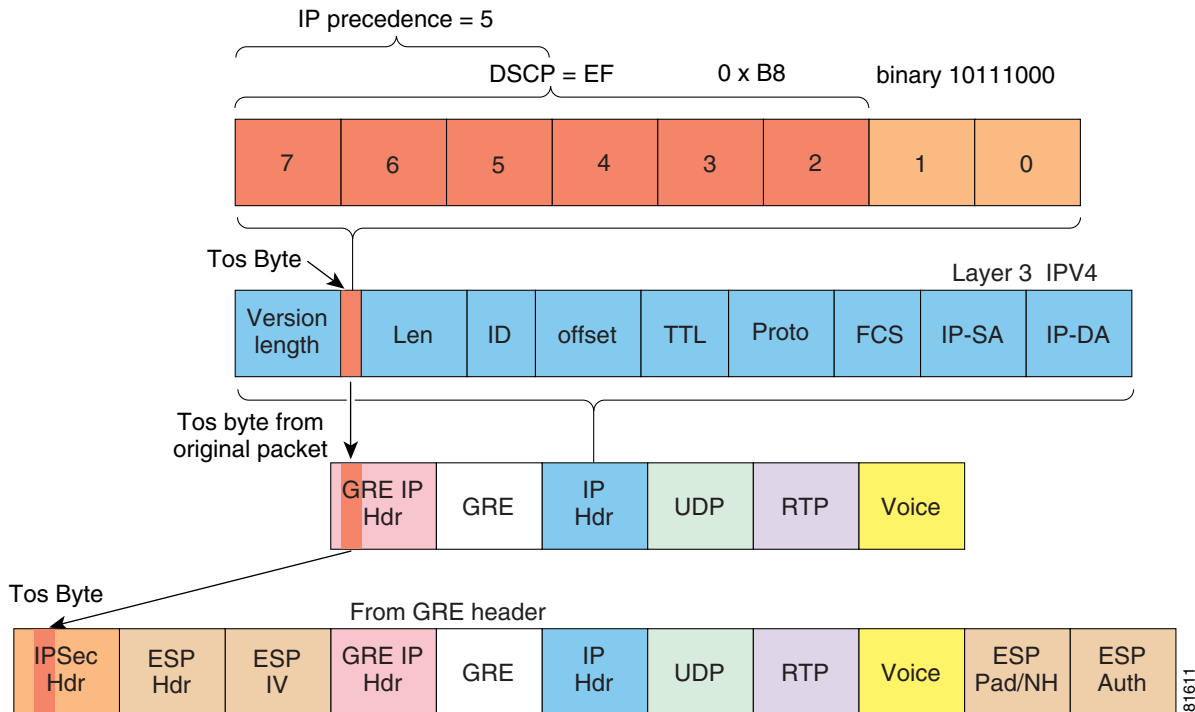
ToS Byte Preservation

In a typical IP Telephony deployment, QoS would be enabled, classifying and marking higher priority traffic (such as voice and H.323 signaling). The ToS byte is commonly used for this purpose.

When encrypting a voice stream the Service Policy can no longer make decisions on the original IP header information, because the original IP header is now encrypted.

However, built into the IPsec protocol standard is the ability to preserve the ToS byte information from the original IP header by automatically copying it to the IP header added by IPsec, so the information is still available for use by Service Policies. Figure 4-6 illustrates this process.

Figure 4-6 IPsec Preserves the ToS Byte



IPsec—RFC 2401 specifies the ToS byte must be copied from the inner header to the outer header. See section 5.1.2.1 *IPv4—Header Construction for Tunnel Mode* in the following document:

<http://www.ietf.org/rfc/rfc2401.txt>

Similarly, with IP GRE, as of Cisco IOS software release 11.3, the ToS byte from the original IP header is copied to the IP GRE header. See the following reference for more information regarding IP GRE:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113t/113t_4/greqos.htm

Also see the “QoS Pre-Classify” section on page 4-12 for more information regarding the interaction of IPsec and QoS Service Policies.

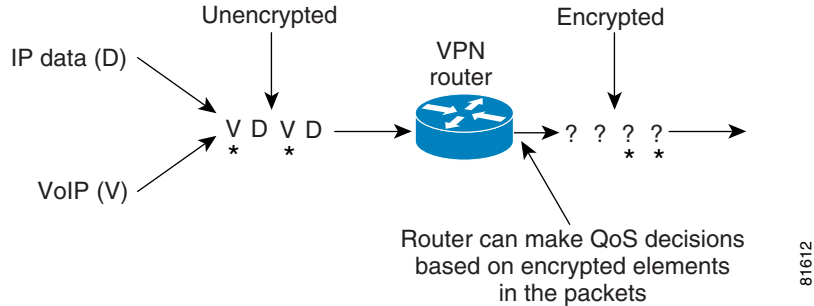
QoS Pre-Classify

QoS Pre-Classify is often confused with the preservation of the ToS byte (see the “ToS Byte Preservation” section on page 4-11) during the packet encryption process. IPsec (and IP GRE) preserve the ToS byte automatically.

In Cisco AVVID solutions, the IP Phone and gateways provide the capability to set the ToS byte so routers can make the appropriate QoS decision. However, most data applications do not set the ToS byte and queuing decisions must be based on other fields of the IP header, including source/destination IP address, port numbers, and protocol.

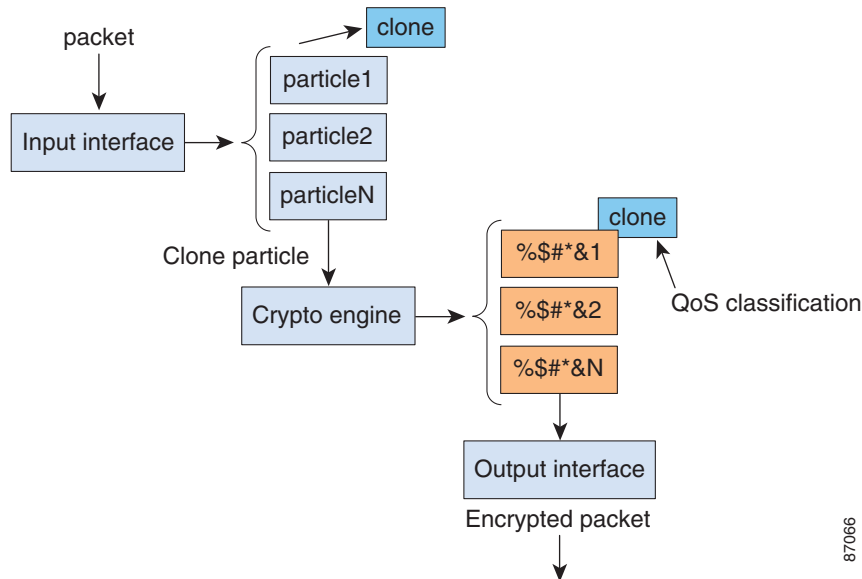
Once the original IP packet is encrypted by IPsec, fields other than ToS byte, such as port numbers, protocol and source/destination IP address fields, are no longer in clear text and cannot match an output service policy. QoS Pre-Classify is a Cisco IOS software feature to allow *fancy queuing*, CBWFQ/WFQ, at the output interface to match on these other fields in the original IP header, even after the original IP header is encrypted. Figure 4-7 illustrates this concept at a high level.

Figure 4-7 QoS Pre-Classify Feature



A key point to remember regarding QoS Pre-Classify is that it is only applicable at the encrypting router's output interface. The fields preserved by QoS Pre-Classify are not available to router's beyond the encrypting router. A simplified illustration is shown in Figure 4-8.

Figure 4-8 How QoS Pre-Classify Works



The following stages summarize the QoS Pre-Classify process illustrated in Figure 4-8:

1. A packet enters the input interface and is stored in particles in a given pool.
2. When the packet matches a crypto map on the interface it would be switched out of, it is passed to the Crypto Engine
3. As part of the Crypto Engine's TX ring processing, a *clone* of the particle, including the IP Header, is created and associated with the packet's data structure.
4. The Crypto Engine encrypts the original packet and places the cipher text in new particle(s), associating the *clone* particle with the new particle(s).

5. If the output interface is not congested, the encrypted packet is simply transmitted.

If the output interface is congested and must be queued for QoS features, the classification can act on the *clone* particle (still in clear text) to match on protocol, source/destination IP address, and port numbers.

Cisco recommends enabling QoS Pre-Classify on all platforms. See the [“Configuring QoS Pre-Classify” section on page 6-23](#) for more information regarding configuration of QoS Pre-Classify.

IP Security (IPSec)

This section outlines the IPSec design options for consideration. It discusses *tunnel* options, and in this design guide, a *tunnel* is specified as an IP GRE tunnel, or an IPSec tunnel. While this document does not go in depth to firewall placement, an example of securing a branch router with access-lists is provided. There is also a discussion on anti-replay and crypto engine QoS.

IPSec and GRE Tunnel Design Considerations

There are currently three recommended design options for a site-to-site IPSec VPN:

- IPSec Tunnel mode—no IP GRE tunnel
- IPSec Transport mode encrypting an IP GRE tunnel
- IPSec Tunnel mode encrypting an IP GRE tunnel (primary recommendation)

This design guide implements IP Tunnel mode encrypting an IP GRE tunnel. The advantages, disadvantages and features and limitations of these options follow.

- **IPSec Tunnel mode—no IP GRE tunnel.** This option does not utilize a IP GRE tunnel. IPSec encrypts IP unicast traffic only, IP Multicast traffic cannot be transported between the IPSec peers without configuring an IP GRE tunnel. This configuration might be sufficient to support the application requirements and its advantage lies in less CPU overhead (primarily at the head-end router) to maintain a IP GRE tunnel to each remote location and a routing protocol’s hello and update packets. IPSec security associations are created for each access list line matched. An access list must be specified in the crypto map to designate packets that are to be encrypted. The access list (when encrypting an IP GRE tunnel) is only one line, a match on protocol 47 (GRE) and the source and destination IP address of the GRE endpoints. When not encrypting a GRE tunnel, it is possible to create an access list which has multiple lines, matching on various portions of the five tuples, source/destination IP address, protocol, source/destination port numbers. A separate security association is created for each access list line match. Each security association has its own ESP (or AH) sequence number. Anti-replay drops can be eliminated or minimized by constructing access lists that create a separate security association for each class of traffic being influenced by per-hop QoS policy.

The *Pre-fragmentation for IPSec VPNs* feature is supported in IPSec Tunnel mode – no IP GRE tunnel, it is first available in Cisco IOS software release 12.1(11)E and is targeted for 12.2(12)T.

- **IPSec Transport mode encrypting an IP GRE tunnel.** This option is commonly implemented; for a G.729 packet it saves 16 bytes per packet over IP GRE tunnels with IPSec Tunnel mode, as an additional IPSec IP header is not required. This byte count would normally be expected to be 20 bytes, the length of an IP header, but 16 bytes as verified by a protocol analyzer, the 4 byte delta would be explained by a different padding length. The IPSec peer IP addresses and the IP GRE peer address must match for transport mode to be negotiated, if they do not match, tunnel mode is

negotiated. The *Pre-fragmentation for IPSec VPNs* feature is **not** supported for Transport mode, as the decrypting router cannot determine if the fragmentation was done prior to encryption or post-encryption by downstream router between the encrypting and decrypting router.

IPSec Transport mode saves link bandwidth, but it does not provide any reduction in packets per second switched by the router. In most instances, packets per second, not packet size, is the limiting factor of a router's main CPU performance.

IPSec Tunnel mode is the default configuration option. To configure Transport mode, it must be specified under the transform set:

```
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
 mode transport
!
```


Configuring IPSec Transport mode to encrypt an IP GRE tunnel provides all the advantages of using IP GRE—it supports IP Multicast, routing protocols and multi-protocol support.

- **IPSec Tunnel mode encrypting an IP GRE tunnel.** This option is implemented in this design guide and in the associated lab testing. It incurs the greatest header overhead of the three options, but it is capable of supporting IP Multicast and the ability to run a dynamic routing protocol within the IP GRE tunnel for failover to an alternative path. It supports Pre-fragmentation for IPSec VPNs. This option was selected for lab testing as it provides the greatest features and flexibility as well as the worst-case scenario in our performance testing in regards to bandwidth consumption.

When configured with a routing protocol running within an IP GRE tunnel, the routing protocol's hello packets maintain the security associations between both (assuming a redundant configuration) head-end routers. There is no need to create a security association to a back-up head-end peer upon failure of the primary peer. Also, routing protocol hello timers (5 seconds by default for EIGRP) can be tuned lower than the hello interval of Internet Security Association and Key Management Protocol (ISAKMP) keepalives—the minimum value is 10 seconds. Detection of a failed head-end peer is quicker when using a routing protocol versus `crypto isakmp keepalive 10`—the dead interval for ISAKMP keepalive is 3 times the keepalive value, or 30 seconds. EIGRP has a default dead interval of three times the hello value of 5 seconds, or 15 seconds.

The diagram in [Figure 4-9](#) illustrates the difference in packet size between IPSec Transport and Tunnel modes.

Figure 4-9 IPSec Transport vs. Tunnel Mode for G.729 Packets



Each option discussed has merit, there is no one design option that is superior to the other alternatives. While this design guide implements IP Tunnel mode encrypting an IP GRE tunnel, internal Cisco deployments of voice over an IPSec Tunnel with no IP GRE tunnel have also been successfully implemented.

Firewall Considerations for Transport of VoIP

Firewall placement at the head-end site depends on the enterprise's security policies. Placing a firewall between the head-end crypto/IP GRE tunnel termination routers and the enterprise network allows the firewall security administrator visibility to the specific port and protocols, as the traffic is unencrypted at that point in the topology.

Placing a firewall between the remote routers and the head-end crypto/IP GRE tunnel termination routers prevents visibility to the specific applications because all traffic is encrypted. IPSec ESP (protocol 50) and UDP port 500 for ISAKMP must be permitted and are the only packets visible to the firewall.

The Cisco IOS Firewall feature set was not tested as part of the Cisco Enterprise Solutions Engineering lab verification of this design as it applies to a split tunneling configuration.

Since this design routes all traffic to the head-end, access to the remote routers from the WAN can be limited by inbound access-lists on the serial interfaces to permit only ISAKMP (UDP port 500) and IPSec ESP (protocol 50) and specific access from the from the head-end routers for management purposes. In this example ICMP is permitted from the upstream router's serial interface but all other ICMP and IP accesses is denied.

```

!
crypto map GRE local-address Loopback0
crypto map GRE 30 ipsec-isakmp
  set peer 192.168.3.1
...
!
interface Loopback0
 ip address 192.168.2.1 255.255.255.252
!
interface Tunnel0
 ip address 10.0.96.1 255.255.255.0
 qos pre-classify
 tunnel source Loopback0
 tunnel destination 192.168.3.1
 crypto map GRE
!
interface Serial0/0.100 point-to-point
 ip address 192.168.1.1 255.255.255.252
 ip access-group 2699 in
 frame-relay interface-dlci 100
  class ts-branch
 crypto map GRE
!
access-list 2699 permit udp host 192.168.3.1 eq isakmp host 192.168.2.1 eq isakmp
access-list 2699 permit esp host 192.168.3.1 host 192.168.2.1
access-list 2699 permit icmp 192.168.1.0 0.0.0.255 192.168.1.0 0.0.0.255
access-list 2699 deny ip any any log-input
!

```

This is for illustrative purposes only. In the above example, console access would need to be provided by a modem for remote management of the device in the event of a failure, as telnet would be denied to the serial interface's IP address.

Anti-Replay Considerations

IPSec offers message integrity, providing for a means to identify if an individual packet is being replayed at a later time. This concept is called connectionless integrity. It also provides for a partial sequence integrity, preventing the arrival of duplicate packets.

These concepts are outlined in RFC 2401 at <ftp://ftp.isi.edu/in-notes/rfc2401.txt>

When ESP Authentication (esp-sha-hmac) is configured in an IPSEC transform set, for each security association, the receiving IPsec peer verifies that packets are received only once. Because two IPsec peers can send millions of packets, a 64-packet sliding window is implemented to bound the amount of memory required to tally the receipt of a peer's packets. Packets can arrive out of order, but they must be received within the scope of the window to be accepted. If they arrive too late (outside of the window), they are dropped.

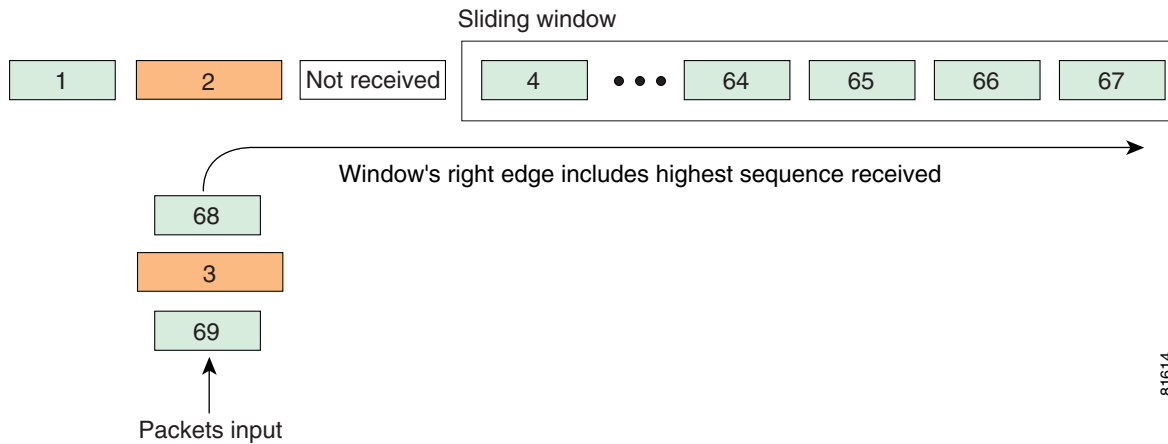
The operation of the anti-replay window protocol is as follows:

1. The sender assigns a unique sequence number (per security association) to encrypted packets.
2. The receiver maintains a 64-packet sliding window, the right edge of which includes the highest sequence number received. In addition, a boolean variable is maintained to indicate if each packet in the current window was received or not.
3. The receiver evaluates the received packet's sequence number:
 - If a received packet's sequence number falls within the window and it has not been previously received, the packet is accepted and marked as received.
 - If the received packet's sequence number falls within the window and was previously received, the packet is dropped and the replay error counter is incremented.
 - If the received packet's sequence number is greater than the highest sequence in the window, the packet is accepted, marked as received, the sliding window is moved "to the right."
 - If the received packet's sequence number is less than the lowest sequence in the window, the packet is dropped and the replay error counter is incremented.

In a converged network implementation with QoS enabled, lower priority packets are delayed such that higher priority packets receive preferential treatment. This has the side effect of also reordering the packets to be out-of-sequence from an IPsec sequence number perspective. Therefore, there is a concern that through the normal QoS prioritization process, packets will be dropped by the receiver as replay errors, when in fact they are legitimately sent/received packets.

Figure 4-10 a visualization of the process. In this example, voice packets 4 through 67 have been received, our data packet "3" was delayed by the queuing process and was transmitted following voice packet "68". When the anti-replay logic is called to process packet "3", it will be dropped, since it will be outside the left edge of the sliding window. Packets might be received out of order, but they must fall within the window to be accepted.

Figure 4-10 Anti-replay Operation



In a converged network of voice and data, anti-replay drops impact data packets rather than voice packets, the QoS configuration prioritizes voice over data. Voice quality should not be impacted by anti-replay.

Anti-replay drops can be eliminated in a pure IPSec configuration (no GRE) by creating separate security associations for voice and data; voice and data packets must match a separate line in the access-list referenced by the crypto map. This is easily implemented if the IP Phones are addressed by a separate network address (RFC 1918 addresses) than the workstations.

Consider the effect of packet loss on a TCP based application. TCP is connection oriented and incorporates a flow control mechanism. The TCP application has no visibility to the reason a packet was dropped. A packet lost by a service policy on an output interface is not different to the application than a packet lost by an anti-replay drop. From a network perspective it would be more efficient to drop the packet before sending it over the WAN link, but the location or nature of the packet loss is immaterial to the TCP driver. Anti-replay drops can be readily created in a lab environment by applying a QoS policy to an output interface and using a traffic generation tool to persistently congest the link with connectionless traffic.

This traffic profile, however, is not representative of most production networks. Table 4-3 represents the traffic profile of a major hospitality organization's core network during a mid-afternoon weekday. NetFlow was enabled on core routers between the organization's branch locations (hotels) and their data center. Voice is not enabled on this network. The primary function of the network is to provide reservation information between the remote hotel and the headquarters data center. DLSw, TN3270, and in-house developed client/server applications are used to access the reservation data on the mainframe. Email and Web applications are also in use.

Table 4-3 Traffic Profile for Major Hospitality Organization

Protocol	Percent Bytes To Branch	Percent Packets To Branch	Percent Bytes To Head-end	Percent Packets To Head-end
ICMP	1.2	1.4	2.2	1.1
UDP	5.6	7.9	10.8	8.1
TCP	93.2	90.7	87.0	90.8

From Table 4-3, TCP is the predominate protocol on this network. The average number of bytes per packet to the branch was 332 bytes and the average bytes per packet to the data center was 159. The traffic profile implemented in testing this design guide is similar to this network, with the addition of approximately 33 percent voice traffic in the total profile. Adding VoIP increases the percentage of UDP packets in the profile, however the data portion continues to be predominately TCP based.

During testing by Cisco's Enterprise Solutions Engineering lab, using the traffic provide documented in this design guide, voice traffic was not adversely affected by anti-replay drops and data drops were typically less than one percent of the total packets decrypted by the receiving router. This drop rate was determined to not adversely impact the function of the network.

Output drops on the output WAN interface tend to be few, if any, and certainly far less than dropped by anti-replay. Anti-replay triggers packet drops more aggressively than the output service policy. This relates to the default size of the output queues and the number of defined classes. In the sample output service policy below, note that each bandwidth class and the class-default can queue a maximum of 64 packets.

```
vpn18-2600-2#show policy-map
  Policy Map llq-branch
    Class call-setup
      Weighted Fair Queueing
        Bandwidth 5 (%) Max Threshold 64 (packets)
    Class mission-critical
      Weighted Fair Queueing
        Bandwidth 22 (%) Max Threshold 64 (packets)
    Class voice
      Weighted Fair Queueing
        Strict Priority
        Bandwidth 168 (kbps) Burst 4200 (Bytes)
    Class class-default
      Weighted Fair Queueing
        Flow based Fair Queueing
        Bandwidth 0 (kbps) Max Threshold 64 (packets)
```

However, the receiving IPSec peer has a single 64-packet anti-replay window (per IPSec Security Association), with which to process all packets from the above priority class (voice), bandwidth classes (call-setup, mission-critical, and internetnetwork-control), and the default class (class-default).

So it stands to reason, anti-replay will be more aggressive than the service policy at dropping packets delayed by voice—due to the size mismatch of the queue depth on output verses the width of the anti-replay window. As more bandwidth classes defined in the policy map, this mismatch increases.

By reducing the queue-limit (Max Threshold) of the bandwidth classes the output service policy becomes more aggressive at dropping packets rather than buffering/delaying them—this further reduces the number of anti-replay drops. The default value of 64 packets is designed to absorb bursts of data traffic and delay, rather than drop, those packets. This is optimal behavior in a non-IPSec enabled network.

With IPSec authentication configured (esp-sha-hmac) in the network, decreasing the queue-limit for the bandwidth classes further reduces anti-replay drops from the 0.5-to-1.5 percent range to tenths of a percent. These configuration changes increase the number of packet drops by the sender's output service policy.

In lab testing using NetIQ Chariot™ with voice plus predominately TCP based data traffic profile, decreasing the queue-limit from 64 to 16 for the critical data traffic, and to 6 for the class-default traffic has shown to decrease anti-replay drops to the tenths of a percent range, while of course, increasing drops on the output interface. The **policy-map** command configuration shown below can be used as a starting point and then tuned by the network manager. Decreasing the queue-limit causes the service policy to be more aggressive in dropping rather than delaying packets and decreases the number of anti-replay drops.

**Note**

The queue limits shown in the following example were tuned for a specific traffic profile and a specific link speed (512 Kbps). They might not apply for different traffic profiles or link speed.

```
policy-map llq-branch
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
    queue-limit 16
  class voice
    priority 168
  class class-default
    fair-queue
    queue-limit 6
```

As a rule of thumb, the **queue-limit** for **class-default** was set lower than **mission-critical**, since the **mission-critical** class contains data traffic which is more important to the enterprise than **class-default**, which is *best effort* traffic in our profile. In this configuration IP Precedence 6 traffic is included in **mission-critical**. If IP Precedence 6 traffic is separated into its own class (**internetwork-control**) use a queue limit of 16 as a starting value.

**Note**

In most networks, the default **queue-limit** command settings and IPSec anti-replay performance is acceptable. Only in situations where there is a requirement to further reduce the affects of IPSec anti-replay and QoS interactions should **queue-limit** tuning be considered. A modification of these values can have other side effects on the QoS service policy and related performance.


Crypto Engine QoS

This section discusses current and future capabilities of hardware crypto engines as they relate to their interaction with QoS.

Current VoIP over IPSec Crypto Engine Capabilities

A crypto engine within a Cisco VPN router's chassis can be viewed as an internal interface that processes packets for encryption or decryption. In Cisco IOS software releases before 12.2(13)T, the crypto engine operates with a FIFO input queue. Packets received from a serial interface for decryption, are interspersed with packets received from an Ethernet interface, to be encrypted. [Figure 4-11](#) illustrates that is no distinction between a voice packet and data packet.

Figure 4-11 FIFO Crypto Engine Illustration



Consider a Cisco 2651XM router deployed at a branch site. It is configured with a full-duplex Fast Ethernet interface, a Serial E1 interface (also full-duplex) and an AIM-BP encryption accelerator. The Fast Ethernet interface connects to the branch's LAN and the Serial interface to the Internet. Consider the factors limiting the throughput of this configuration:

- Clock rate of the slowest interface—E1 rate transmitted and received (approximately 4 Mbps)

- Packet forwarding rate of the router's main CPU—in packets per second
- Crypto engine encryption/decryption rate—in packets per second

The performance characteristics of the above items, are further influenced by the traffic mix—including the size of the IP packets being switched through the network, the switching path (process, fast, CEF) of the packets and the features present in the configuration. In most hardware platforms, the packet per second capabilities of the router are more important for planning purposes than bits per second switched through the router. If the average packet size switched through the router increases from 128 bytes to 256 bytes, the packet per second capabilities of the main CPU is not necessarily cut in half.

Additionally, the control plane requirements of the internetwork—the number of routes in the routing table, the overall network stability and requirements of the routing protocol in use, the network management (SNMP) requirements, additional features enabled on the router—DLSw, TACACS, NTP, QoS, access control lists, all consume CPU resources. Assume that the amount of available memory for the main CPU, interfaces and crypto engine are adequate, no memory limitations exist.

The ratio of packets switched through, and originated by, the router in relation to those selected by the crypto map's access-list for encryption/decryption must also be considered. If encrypting an IP GRE tunnel, this tends to be a large percentage of encrypted to total packets. If not encrypting a IP GRE tunnel, and selecting only a portion of the total data traffic from the LAN/WAN interface, the ratio could be quite small.

The hardware crypto engine accelerator becomes congested when its packet processing capabilities are less than those of the router's main CPU and interface clock speed.

In cases where congestion occurs in the crypto engine, it is possible for the crypto engine to become over-subscribed—either on a momentary or sustained basis. In these cases, there are three possibilities:

- The over-subscription does not affect VoIP packets to the extent that voice quality issues result.
- Additional processing latency occurs, adding some unnecessary delay and/or jitter into voice streams.
- Voice streams experience some degree of packet loss—affecting voice quality.

Cisco internal testing and evaluation has shown it to be extremely difficult for conditions to arise that cause crypto engine congestion. In nearly all cases, the Cisco VPN Router platform's main CPU is exhausted prior to reaching the limit of the crypto engine's packet processing capabilities.


Nevertheless, Cisco provides a solution to the *potential* problem so that networks do not encounter such a situation. Consequently, Cisco developed the *LLQ for Crypto Engine* feature within Cisco IOS software—explained in the [“LLQ for Crypto Engine”](#).

LLQ for Crypto Engine

Starting with Cisco IOS software release 12.2(13)T, the crypto engine has been enhanced to provide LLQ for the Crypto Engine. This entails providing a dual-input queuing strategy: a priority or Low Latency Queue; and a best effort queue. The feature is targeted at alleviating any effects of momentary or sustained over-subscription of the hardware crypto engine, which can result in priority traffic (such as voice and video) experiencing quality issues.

The classification component to segregate traffic between the priority (LLQ) and the best effort queue is based on the CBWFQ service policy on the output interface(s). See [Figure 4-12](#).

Figure 4-12 LLQ Crypto Engine Queue Illustration



There is no additional configuration required to enable LLQ for Crypto Engine; it is enabled by the presence of a CBWFQ Service Policy on an output interface of the VPN router.

Traffic specified in the CBWFQ service policy to be included in the priority queue(s) (LLQ) will be sent to the Crypto Engine's LLQ. Traffic included in any bandwidth classes (queues) and default class queue will be put in the Crypto Engine's best effort queue.

It is possible for more than one output interface to be configured, each potentially having a CBWFQ service policy. However, the Crypto Engine acts like a single *interface* inside the VPN router, encrypting/decrypting all outbound/inbound traffic streams for each interface on which Crypto is applied. In the case of multiple CBWFQ service policies (on different interfaces), the Crypto Engine maps all priority queues (LLQ) to its LLQ and all other queues to its best effort queue.

The priority (LLQ) queue for the crypto engine is similar in function to the priority (LLQ) queue for a service policy attached to an output interface. The Crypto Engine is analogous to an interface from Cisco IOS software's perspective.

Further, in the event that the Crypto Engine becomes oversubscribed—for short durations or sustained periods—the LLQ for Crypto Engine feature insures that if packets are dropped by the Crypto Engine, they are of appropriately low priority (not VoIP packets).

Although the feature is enabled by the presence of a CBWFQ Service Policy, like QoS it does not actually engage prioritization via the two-queuing strategy until the crypto engine itself experiences congestion.

As software-based crypto adds unacceptable latency and jitter, there are no plans to incorporate this feature for software crypto. This design guide recommends hardware acceleration of IPSec for V³PN deployments.

When is LLQ for Crypto Engine Required

The *LLQ for Crypto Engine* feature in Cisco IOS software is not a prerequisite for deploying many V³PN implementations in a high quality manner. As indicated previously, internal Cisco evaluations found it extremely difficult to produce network traffic conditions that resulted in VoIP quality suffering.

However, the feature should be viewed in the same light as Cisco QoS: it is there in Cisco IOS software to safeguard against degradation of high priority traffic delivery in periods of *harsh* network conditions when momentary or sustained over-subscription can occur.

In general, the *LLQ for Crypto Engine* feature offers the most benefit under one of the following conditions:

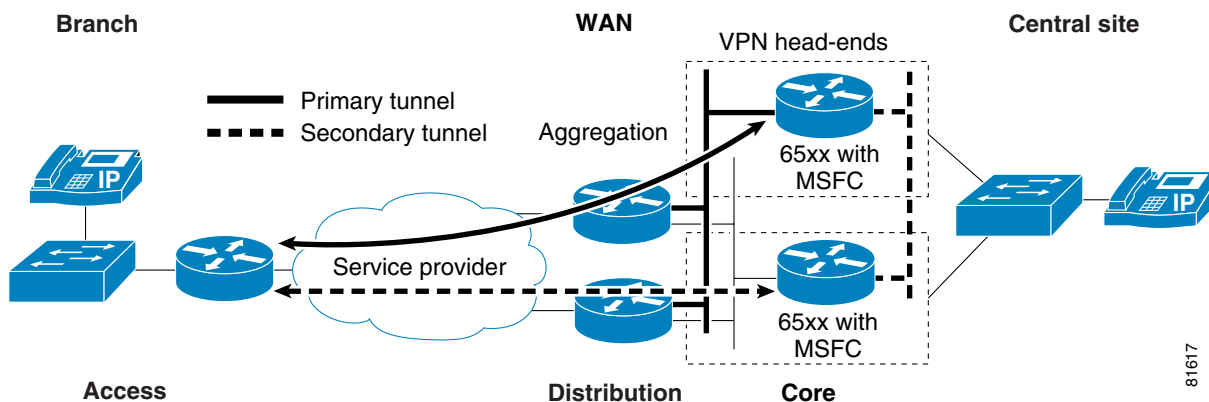
- When implementing Cisco IOS VPN Router platforms that have a relatively high amount of main CPU resources relative to Crypto Engine resources.
- When the network experiences a periodic or sustained *burst* of large packets (for example, video applications).

To summarize, high-quality V³PN deployments are possible today without the *LLQ for Crypto Engine* feature in Cisco IOS software. The addition of this feature in Cisco IOS software further insures that high priority applications such as voice and video can operate in a high-quality manner even under harsh network conditions.

Head-end Topology

This design can be considered an alternative to a typical private WAN deployment where a hub and spoke topology is deployed which transports voice traffic. As such, it is assumed the rules of scalable and redundant network design need to be present in the solution. This Design Guide recommends incorporating the IPsec/GRE head-end devices into a scalable hierarchical network model. The hierarchical network design model is represented in three layers: *core*, *distribution* and *access*. Each layer provides a different functionality. Figure 4-13 illustrates how this design overlays the model presented in this publication.

Figure 4-13 Solution Deployment Topology



In the hub and spoke IPsec VPN topology, the Branch routers represent the *access* layer. These routers terminate the WAN interface(s) to the service provider, connect to the local LAN segment, and provide IPsec and GRE tunnel termination. Increasing availability also increases costs, but ideally multiple serial interfaces with PVCs to each WAN Aggregation router would be recommended.

The *distribution* layer is implemented with pairs of WAN Aggregation routers. In an implementation over a Layer 3 service provider (Internet Service Provider), these routers would have a high-speed WAN interface(s) and one or more Fast Ethernet or Gigabit Ethernet interfaces. They would typically be eBGP peers with the ISP's edge routers and would peer via iBGP between all WAN Aggregation routers. They would also implement QoS on the WAN interface to prioritize the voice traffic.

In an implementation with a Frame Relay service provider or point-to-point network, the *distribution* layer WAN Aggregation routers would have one or more high-speed WAN interfaces and perhaps hundreds of subinterfaces—one for each branch.

The *core* layer is implemented with pairs of Layer3/Layer2 switches, ideally Catalyst 6500 with MSFC. These switches provide connectivity to the network core and provide redundant Layer 2 or Layer 3 connectivity between the WAN Aggregation routers and the IPsec/GRE head-end routers.

The IPsec/GRE head-end routers are Cisco 7200VXR or Cisco 3600/3700 series routers. Their LAN connectivity is provided by the Catalyst 6500 switches and is typically implemented as dual Fast Ethernet attached routers. They terminate the IPsec peers as well as the IP GRE tunnels. They advertise the branch subnets learned through the tunnel interfaces to the core switches. Each branch router provides two IP GRE tunnels, one to each pair of IPsec/GRE head-end routers. This provides for an

alternate path in the event a head-end router is taken out of service. There is no need for QoS configured on these routers, since the LAN interfaces will not be congested. The main CPU must be monitored and managed as part of the enterprise's capacity planning function.

The design recommendation of terminating the IPSec/GRE traffic on routers dedicated for this purpose is consistent with design principles proven by large scale deployments of SNAsw, TN3270 Server or DLSw. Dedicating routers for a specific function at the network core provides several advantages. The performance characteristics of the IPSec/GRE head-end might be dramatically different than a WAN aggregation router. Separating the two functions allows different ratios between IPSec/GRE head-ends and WAN aggregation: two WAN aggregation routers might be sufficient for four IPSec/GRE head-ends.

WAN aggregation routers might require a lower average CPU busy percentage to accommodate CPU spikes. WAN aggregation routers running BGP might experience considerable CPU spikes during network instability due to link flaps, for example. Also, the separation allows different versions of Cisco IOS software to be run at different locations in the network topology. For example, the WAN aggregation router might be running a General Deployment (GD) release of 12.0 mainline, while the IPSec/GRE head-end routers need an ED (Early Deployment) release to support a new hardware encryption accelerator.

The advantages and flexibility of this design outweigh the additional costs involved by separating the WAN aggregation from the IPSec/GRE head-end.

Head-end Router Locations

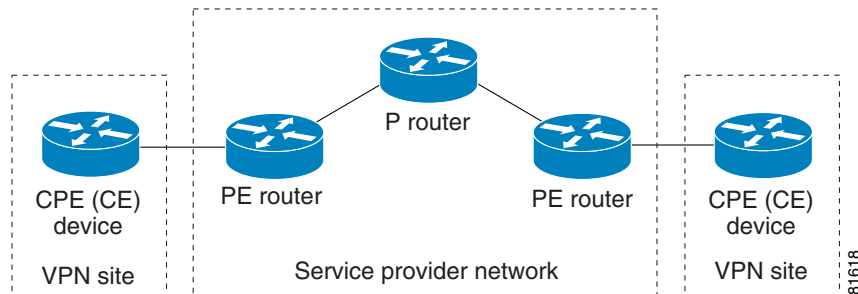
The number and geographic locations of the head-end routers require careful consideration. If using a Layer 3 service provider (Internet Service Provider), consider the geographic location of the majority of users/sites and the placement of head-end routers. For example, Cisco has a large concentration of employees in San Jose, California and Research Triangle Park (RTP), North Carolina. It's practical to locate head-end routers at both locations, with West coast sites terminating in San Jose and East coast locations at RTP. Dual head-end routers would be installed at each location, physically diverse within the campus. For global corporations, head-end routers in EMEA and Asia-Pacific should also be considered. A large portion of the voice delay budget will be consumed by the service provider, therefore the goal is to minimize the time spent in the service provider's cloud.

Service Provider Recommendations

For a V³PN deployment to operate successfully, the enterprise network designer must address the QoS requirements of encrypted voice traffic if Layer 3 service providers transport traffic between branch and head-end devices.

Boundary Considerations

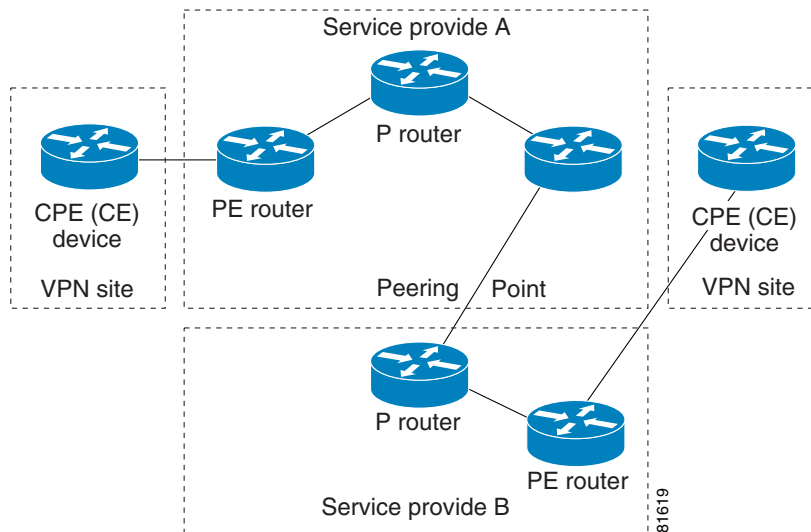
In a Layer-3 service provider deployment, the Enterprise organization and the service provider must have identical QoS policies implemented on the link between the CPE (Customer Premise Equipment or CE Customer Edge) device and the PE (Provider Edge) router. [Figure 4-14](#) illustrates this simplified topology.

Figure 4-14 Service Provider/Enterprise Boundary

Additionally, for the service provider to offer a QoS enabled network, there needs to be a control point at the PE routers to either limit or police the amount of high priority traffic (voice, based on ToS byte) or provide some accounting function to bill based on the priority level of the traffic. Without this control in place, there is no incentive for the organization to accurately mark packets based on their required priority—they could mark all traffic as highest priority, intentionally or inadvertently.

Cross-Service-Provider Boundaries

If multiple Layer 3 service providers are used to connect the enterprise's sites, the complexity increases. The enterprise must coordinate the CPE to PE QoS with two separate ISP's. This adds time and complexity to the implementation. See [Figure 4-15](#).

Figure 4-15 Multiple Service Providers

To have end-to-end QoS across the IPSec VPN, both service providers must agree on the amount of high priority traffic to be accepted between them. Many ISPs set the IP Precedence/DSCP value to 0 at their peering points and PE routers.

Implementations spanning multiple ISPs are more difficult to implement and manage. A contiguous service provider implementation is recommended, but when not possible, an enterprise is encouraged to seek an agreement with service providers specifying how high priority traffic across boundaries will be handled.

Service Level Agreements (SLA)

In order to support V³PN, Cisco is offering a new Cisco Powered Network Designation called *IP Multi-service VPN* that qualified service providers can attain, signifying that they are capable of offering such VPN services.

Per the Cisco Powered Network Service Provider requirements, they must meet these minimum SLA components:

- Jitter—Less than or equal **20 msec**
- Delay—Less than or equal **60 msec one way**
- Packet Loss—Less than or equal **0.5 percent**

They are responsible for meeting the terms of the SLA they provide to the enterprise organization, just as they would if providing service via a Private WAN service offering, such as Frame Relay or ATM.

Cisco Powered Network References

These Cisco documents are useful references for V³PN implementations:

- *IP Multi-service VPN Additional Requirements Summary*—<http://www.cisco.com/warp/public/779/servpro/cpn/join/criteria.html>
- Search for designated Cisco Powered Network service providers—http://www.cisco.com/cgi-bin/cpn/cpn_pub_bassrch.pl



Note Please select “VPN/IP–Multi-service” in the selection list on the search page.

- Network Service Definitions—<http://www.cisco.com/warp/public/779/servpro/cpn/glossary.html>
- Cisco Powered Network program membership application—http://www.cisco.com/cgi-bin/cpn/cpn_screen_zero.pl

Enterprise organizations can use these links to locate Cisco Powered Network service providers, and service providers can use them to enroll in the program.

Load Sharing

To accommodate higher traffic volumes, some implementations might require multiple physical links—such as two T1 links—for a large branch office. This section addresses several load-sharing topics:

- [Load Sharing Capabilities, page 4-27](#)
- [Encrypted Traffic Appears as a Few, Large Flows, page 4-27](#)
- [Minimize Out-of-Order Packets, page 4-27](#)
- [Load Sharing Design Approach, page 4-28](#)
- [Load Sharing from Head-end to Branch, page 4-30](#)
- [Service Provider Considerations for Load Sharing, page 4-32](#)

Load Sharing Capabilities

Routing protocols—such as OSPF, EIGRP, or BGP—have the ability to insert multiple equal cost or (in the case of EIGRP) unequal cost routes into the routing table. Path determination is selected from routes in the routing table rather than what might be contained in the routing protocol's topology database. The Cisco IOS fast-switching path will load share on a per destination basis and CEF by default will load share on a per source/destination basis and can be configured for per-packet load sharing. Since CEF includes the source address in the decision process, it provides for a more granular distribution than fast switching. Packets that are process switched will be per-packet load shared; however, process switching is not recommended as it negatively impacts the overall throughput capabilities of the router.

With fast-switching and CEF-switching, load sharing on links tends to be equally balanced as the number of IP flows (source/destination IP addresses communicating) increase. It is advantageous to transmit all the packets for a flow across a single interface because the likelihood of the packets being received out-of-order is less than if the packets are sent over multiple links. If there are only a few IP flows, and one of these flows is a high bandwidth consumer and the remaining flows are low bandwidth consumers, the amount of link utilization will be far from equal. An example of this would be one workstation sending a large file transfer (FTP) and the remaining stations all using Telnet. The link utilization across two links in this situation is far from equal.

A solution to equally balance the two physical links has been to process switch (not recommended due to the performance constraints) or to use CEF's per-packet load sharing. Another approach is to bundle multiple links at Layer 2 so they appear to the routing protocol to be one link and let the interface driver or Layer-2 logic address load sharing. Examples of this are Inverse Multiplexing over ATM (IMA) and Multilink PPP.

Encrypted Traffic Appears as a Few, Large Flows

IPSec tunnels will *hide* the source and destination address of the traffic selected for encryption in its own IP header and the switching decision at the physical interface will be presented with high bandwidth flows from a few sources—the number of IPSec tunnels. Internet Key Exchange (IKE) and any traffic not selected for encryption by the crypto map (split tunneling or other management traffic) will also be seen on the interface—but this traffic will be assumed to be minimal.

Minimize Out-of-Order Packets

Out of order packets are detrimental from TCP's perspective in that they must be re-ordered by the receiving station's TCP stack before being delivered to the application. This consumes memory and decreases throughput. For voice, some out-of-order packets can be tolerated (MAX_MISORDER defines the maximum mis-order of packets allowed by RTP); however, the packets must be correctly reordered by the jitter buffer of the receiving phone. The human ear will not tolerate listening to digitized voice in anything but the correct order. It will sound garbled or unintelligible.

From IPSec's perspective, per packet load sharing for the same IPSec flow after the packets have been encrypted and assigned an ESP sequence number will increase the probability that packets will be dropped due to replay protection checks. At the same data rate, CEF load sharing per-packet on equal cost IPSec-encrypted GRE tunnels (**ip load-sharing per-packet** command on multiple tunnel interfaces) will show less anti-replay drops than routing all the traffic in one IPSec encrypted GRE tunnel and using CEF load sharing per-packet (**ip load-sharing per-packet** command on the WAN interfaces) to switch the encrypted tunnel's packets out multiple interfaces.

However, reducing the anti-replay drops is not the only goal. Per-packet load sharing across two GRE tunnels increases the likelihood the voice packets between any two IP Phones will take different paths through the network—increasing the likelihood that they will arrive out of order.

Load Sharing Design Approach


To balance these somewhat conflicting requirements, the best approach when implementing multiple physical links would be to CEF switch (per source/destination load sharing, not per-packet) on two equal-cost GRE tunnels and configure each GRE tunnel so it has an affinity to a particular interface when all interfaces are up. This approach strives to maintain voice packets between any two IP phones in the same IPSec/GRE tunnel and on the same physical interface. With voice and data packets using both IPSec/GRE tunnels, and these tunnels each routed over a separate physical interface, both links will be used while maintaining the same path for any one particular flow.

This approach will not provide a precise distribution of packets over the two links as would per-packet load sharing, but it will allow both links to be used and minimizes the negative aspects of out of order packets for voice and data post decryption and anti-replay drops when those packets are in the IPSec tunnel.

In the event one of the serial links fails, the service policy attached to the physical interface must be configured to provision sufficient bandwidth for the number of calls allowed by the sites's Call Admission Control process—CallManager *locations* or use of a Gatekeeper. When both links are up and operational, the site will never consume all the bandwidth provisioned for the voice LLQ, as the Call Admission Control process will limit the number of calls. Bandwidth not used by the LLQ is not wasted, it will be used by the other bandwidth classes and class-default.

Use the topology diagram illustrated in [Figure 4-16](#) as a guide for the examples shown.

Figure 4-16 Load Sharing Topology Example



The router on the left of the diagram represents the branch router and it has two GRE tunnels to the head-end IPSec/GRE peers.

```
vpnjk-2600-18#show run | include tunnel
interface Tunnel0
 tunnel source Loopback0
 tunnel destination xx.0.32.22
interface Tunnel1
 tunnel source Loopback0
 tunnel destination xx.0.32.23

vpnjk-2600-18#show ip interface brief | include Loopback
Loopback0                xx.0.218.1             YES NVRAM up   up

vpnjk-2600-18#show run | include ip route xx.0.32
ip route xx.0.32.22 255.255.255.255 Serial0/0.100
```

```
ip route xx.0.32.23 255.255.255.255 Serial0/0.101
```

To configure an affinity for the logical (IPSec/GRE) links to the physical links, two host static routes are configured. An alternate approach would be to advertise a dynamic route with a more specific mask (in this example using 255.255.255.255 or a host route) down each WAN interface in addition to a summary, supernet, or less specific route down each WAN link. The remote router will route both IPSec/GRE tunnel interfaces out the remaining link in the event of a WAN link failure with this configuration. The routing table in this example looks like this:

```
vpnjk-2600-18#show ip route | begin xx.0.0.0
      xx.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
B       xx.0.0.0/8 [20/0] via xx.0.32.6, 00:16:07
          [20/0] via xx.0.32.2, 00:16:07
S       xx.0.32.23/32 is directly connected, Serial0/0.101
S       xx.0.32.22/32 is directly connected, Serial0/0.100
C       xx.0.32.4/30 is directly connected, Serial0/0.101
C       xx.0.32.0/30 is directly connected, Serial0/0.100
C       xx.0.218.1/32 is directly connected, Loopback0
      10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
D       10.2.0.0/24 [90/297270016] via 10.96.1.2, 00:16:07, Tunnel0
          [90/297270016] via 10.96.1.6, 00:16:07, Tunnel1
```

Note that the remote router is learning a route to xx.0.0.0/8 via BGP and with **maximum-paths 2** configured both are inserted in the routing table. In this example the GRE tunnel's interface *delay* has not been changed and EIGRP inserts both routes to 10.2.0.0/24 in the routing table as they are equal cost. The BGP configuration for the remote router is as follows:

```
router bgp 65018
  no synchronization
  bgp log-neighbor-changes
  network xx.0.218.1 mask 255.255.255.255
  neighbor xx.0.32.2 remote-as 65000
  neighbor xx.0.32.6 remote-as 65000
  maximum-paths 2
  no auto-summary
!
```

CEF is enabled on the branch router and by default will load share per source/destination IP address rather than per-packet:

```
vpnjk-2600-18#show ip cef 10.2.0.0 detail
10.2.0.0/24, version 27, epoch 0, per-destination sharing
```

In this illustration, a traffic generation tool is sending traffic from two different source addresses with three different destination addresses:

```
Summary of IP traffic streams on FastEthernet0/1
  ts#   tos  len  id frag ttl protocol  chksm  source          destination
  ---  ---  ---  --- --- ---  ---  ---  ---  ---
  1 TCP  48   576 0000 0000 60    6    67D1 10.96.0.17      10.2.0.45
  2 UDP  B8    60 0000 0000 60    17   68CF 10.96.0.18      10.2.0.43
  3 UDP  40   188 0000 0000 60    17   6967 10.96.0.17      10.2.0.24
  4 UDP  B8    60 0000 0000 60    17   696E 10.96.0.18      10.2.0.24
```

To verify that both IPSec/GRE tunnel interfaces are being used as well as both physical interfaces, the network manager can use the **show ip cef exact-route** command and either a **show interface** command or in this case, with Frame Relay interfaces, the **show frame pvc** command to verify both physical interfaces are being used.

```

vpnjc-2600-18#show ip cef exact-route 10.96.0.17 10.2.0.45
10.96.0.17      -> 10.2.0.45      : Tunnel0 (next hop 10.96.1.2)
vpnjc-2600-18#show ip cef exact-route 10.96.0.17 10.2.0.24
10.96.0.17      -> 10.2.0.24      : Tunnel0 (next hop 10.96.1.2)
vpnjc-2600-18#show ip cef exact-route 10.96.0.18 10.2.0.43
10.96.0.18      -> 10.2.0.43      : Tunnel0 (next hop 10.96.1.2)
vpnjc-2600-18#show ip cef exact-route 10.96.0.18 10.2.0.24
10.96.0.18      -> 10.2.0.24      : Tunnel1 (next hop 10.96.1.6)

vpnjc-2600-18#show frame pvc 101 | include output pkts
  input pkts 37          output pkts 3171          in bytes 4145
vpnjc-2600-18#show frame pvc 100 | include output pkts
  input pkts 154         output pkts 7158           in bytes 20525

```

The **show adjacency** command can also be used to verify the packet distribution.

```

vpnjc-2600-18#show adjacency serial 0/0.101 detail | include packet
                               33555 packets, 15210652 bytes
vpnjc-2600-18#show adjacency serial 0/0.100 detail | include packet
                               15427 packets, 2170788 bytes
vpnjc-2600-18#show adjacency tunnel 0 detail | include packet
                               16541 packets, 1389444 bytes
vpnjc-2600-18#show adjacency tunnel 1 detail | include packet
                               37140 packets, 14626480 bytes

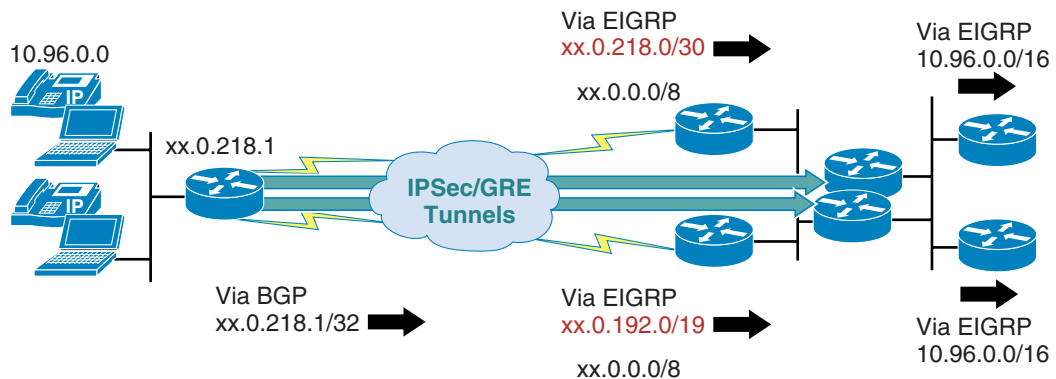
```

Load Sharing from Head-end to Branch

The previous load sharing discussions focused on configuration from the branch router's perspective. If the WAN links are Frame Relay or dedicated T1 links between branch router and head-end, the enterprise can control load sharing for the head-end to branch traffic by configuring the head-end WAN aggregation routers to advertise (or learn dynamically from the branch) the appropriate more specific routes to the IPsec/GRE head-end routers.

Review the sample topology shown in [Figure 4-17](#).

Figure 4-17 Load Sharing from Head-end to Branch



The remote router (on the left of the diagram) will advertise its loopback interface IP address to both head-end WAN routers as xx.0.218.1/32. This loopback address is the IPsec/GRE tunnel destination address for both head-end IPsec/GRE routers. The top WAN aggregation router has a static route for xx.0.0.0/8 and xx.0.218.0/30 to the Null0 interface and is redistributing both these routes to the IPsec/GRE head-end routers via EIGRP 23. The lower WAN aggregation router has a static route for

xx.0.0.0/8 and xx.0.192.0/19 to the Null0 interface and is also redistributing these routes via EIGRP 23. Note that xx.0.218.0/30 and xx.0.218.1/32 are both more specific (have a longer prefixes) which fall in the address space of xx.0.192.0/19. An example of the WAN router's configuration follows.

```
upperWAN#show run | include Null0
ip route xx.0.0.0 255.0.0.0 Null0
ip route xx.0.218.0 255.255.255.252 Null0

upperWAN#sh ip route | include xx.0.218
B      xx.0.218.1/32 [20/0] via xx.0.32.1, 1d02h
S      xx.0.218.0/30 is directly connected, Null0

upperWAN#show run | begin router eigrp 23
router eigrp 23
 redistribute static
 network xx.0.0.0
 default-metric 1000 100 255 1 1500
 ...

lowerWAN#show run | include Null0
ip route xx.0.0.0 255.0.0.0 Null0
ip route xx.0.192.0 255.255.224.0 Null0

lowerWAN#show ip route | begin xx.0.218
B      xx.0.218.1/32 [20/0] via xx.0.32.5, 6d03h
D      xx.0.218.0/30 [90/2588160] via xx.0.1.20, 1d03h, FastEthernet0/1
S      xx.0.192.0/19 is directly connected, Null0

lowerWAN#show run | begin router eigrp 23
router eigrp 23
 redistribute static
 network xx.0.0.0
 default-metric 1000 100 255 1 1500
 .....
```

The lower IPsec/GRE termination router will be configured to ignore any routing advertisements learned via EIGRP 23 for the xx.0.192.0/19 network which have a prefix greater than or equal to 30 bits. The upper IPsec/GRE termination router doesn't have this distribute-list filtering based on the IP Prefix-list feature. IP routing is accomplished by using the matching route that has the longest prefix. The net result of this configuration provides for the upper IPsec/GRE router to follow the xx.0.218.0/30 route advertised by the upper WAN aggregation router while the lower IPsec/GRE router will follow the xx.0.192.0/19 advertised by the lower WAN aggregation router. Since both WAN aggregation routers have a matching route for xx.0.218.1 learned via BGP from the remote branch, they will route the tunnel out their respective WAN interfaces.

In the event one of the WAN routers fail, the IPsec/GRE head-end routers will use as a last resort the xx.0.0.0/8 route advertised by the surviving WAN aggregation router to reach the tunnel destination. If the surviving WAN router is the lower router, the xx.0.192.0/19 will continue to be advertised to both IPsec/GRE routers and that route will be the most specific route to reach xx.0.218.1.

The EIGRP configuration for the lower IPsec/GRE termination router follows:

```
router eigrp 23
 network xx.0.0.0
 distribute-list prefix FOLLOWslash19 in
 ... [additional commands removed] ...
!
router eigrp 45
 network 10.0.0.0
```

```

... [additional commands removed] ...
!
ip prefix-list FOLLOWslash19 seq 5 deny xx.0.192.0/19 ge 30
ip prefix-list FOLLOWslash19 seq 100 permit 0.0.0.0/0 le 32

```

The instance of EIGRP 45 is included for the 10.0.0.0 network which is used in the remote, tunnel and head-end routers. Both IPSec/GRE routers will advertise an equal cost route to the remote subnet, 10.96.0.0/16, so the IP packets from the head-end LANs will use both GRE tunnels to reach the branch subnet. As with the branch configuration, the default of IP CEF per source/destination load sharing will be in effect to send voice and data traffic down each GRE tunnel. Another alternative would be to use multiple HSRP groups on the IPSec/GRE routers so traffic from the head-end arrives on both head-end routers.

Following is an excerpt of the routing table from both IPSec/GRE routers.

```

upperIPSecGRE#show ip route | begin xx.0.0.0
    xx.0.0.0/8 is variably subnetted, 15 subnets, 5 masks
C       xx.0.1.0/24 is directly connected, FastEthernet0/1
D       xx.0.0.0/8 [90/2588160] via xx.0.1.21, 02:07:35, FastEthernet0/1
        [90/2588160] via xx.0.1.20, 02:07:35, FastEthernet0/1
D       xx.0.32.23/32 [90/156160] via xx.0.1.23, 05:11:16, FastEthernet0/1
C       xx.0.32.22/32 is directly connected, Loopback0
... [detail removed] ...

D       xx.0.218.0/30 [90/2588160] via xx.0.1.20, 02:07:35, FastEthernet0/1
D       xx.0.192.0/19 [90/2588160] via xx.0.1.21, 05:09:03, FastEthernet0/1

lowerIPSecGRE#show ip route | begin xx.0.0.0
    xx.0.0.0/8 is variably subnetted, 14 subnets, 5 masks
C       xx.0.1.0/24 is directly connected, FastEthernet0/1
D       xx.0.0.0/8 [90/2588160] via xx.0.1.21, 02:10:20, FastEthernet0/1
        [90/2588160] via xx.0.1.20, 02:10:20, FastEthernet0/1
C       xx.0.32.23/32 is directly connected, Loopback0
D       xx.0.32.22/32 [90/156160] via xx.0.1.22, 02:10:17, FastEthernet0/1

... [detail removed] ...
D       xx.0.192.0/19 [90/2588160] via xx.0.1.21, 05:14:01, FastEthernet0/1

```

With this configuration, provided both WAN routers are available, the IPSec/GRE tunnels will have an affinity to their particular WAN link.

Service Provider Considerations for Load Sharing

If the enterprise is using two Internet service providers to terminate the branch router and head-end, the head-end configuration can influence the path selection. It can cause the head-end to route packets from the one IPSec/GRE head-end out service provider “A” and from the second IPSec/GRE head-end out service provider “B”—while maintaining the logical to physical affinity concept established in the earlier sections.

However, if the physical links terminate to the same service provider—with either on the same router or a different router at the service provider point-of-presence (POP)—then a more-specific (host) route must be used by the service provider. In this case, the enterprise organization must source the branch router’s tunnel interface off a distinct IP address rather than using the same address (Loopback 0) as shown in [Figure 4-17](#). The service provider might not offer this service to the enterprise organization due to the added configuration complexity and additional routes within its network.

Because service provider offerings vary from service provider to service provider, it is advisable to review the requirements of the enterprise network with proposed service providers prior to implementation.

E911 and 911 Emergency Services

Handling of emergency calls would be implemented on the same model as a centralized CallManager with remote IP Phones on a traditional private Frame Relay deployment. Please refer to the following URL for an explanation of addressing emergency calls in a Cisco AVVID deployment:

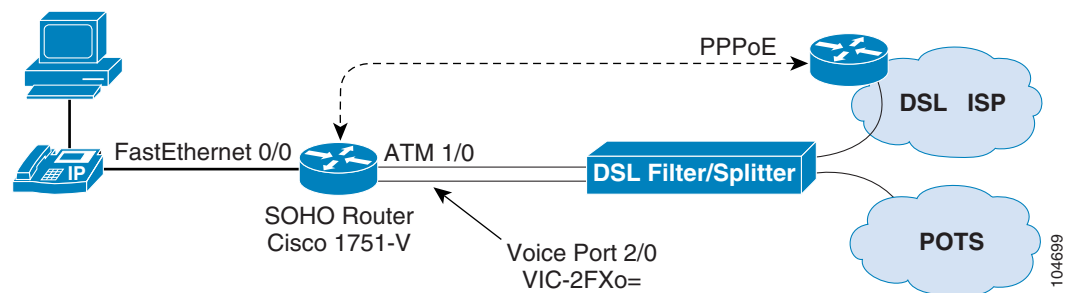
- http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidances_list.html

Survivable Remote Site Telephony

A Survivable Remote Site Telephony (SRST) configuration would be desirable in a V³PN design which utilized a centralized CallManager configuration. In the event the IPSec VPN tunnel is down, remote locations would continue to have limited support and use of their IP Phones. In-depth configuration and testing of SRST is beyond the scope of this document, however proof of concept testing was completed.

The configuration shown in the following diagram uses a Cisco 1751 Virtual Private Network (VPN) Module (MOD1700-VPN), a WIC-ADSL and a VIC-2FXO running Cisco IOS image c1700-k9o3sv8y7-mz.122-13.T1. This router was connected to Cisco and the production CallManager using an ISP and the DSL interface. The analog phone line which is associated with the DSL service was connected to the FXO port of the router. A common DSL filter/splitter was used to split the DSL and analog line to the router. Refer to [Figure 4-17](#).

Figure 4-18 SRST Topology



The following represents the relevant SRST configuration commands for the above topology.

```
!
ip dhcp pool Client
  import all
  network 10.81.2.0 255.255.255.248
  default-router 10.81.2.1
  dns-server 64.102.6.247 171.68.226.120
  domain-name cisco.com
  option 150 ip 64.102.2.93
  netbios-name-server 171.68.235.228 171.68.235.229
!
voice-port 2/0
  connection plar 23685
```

```
description My Home Phone Line
!
voice-port 2/1
!
!
dial-peer voice 45 pots
 destination-pattern 9
 port 2/0
!
call-manager-fallback
 ip source-address 10.81.2.1 port 2000
 max-ephones 2
 max-dn 2
 access-code fxo 9
!
```

The Cisco 1751 router was configured as a DHCP server for the IP Phone, which is a requirement for SRST to function. The phone would have initially needed to register with the production CallManager to function. The phone in question has the five digit extension of 2-3685. The goal of this configuration is to allow the phone to call out the analog POTS line in the event the IPSec tunnel is down and to route incoming calls on the analog POTS line to the IP Phone at extension 2-3685.

To test this configuration, the RJ-11 jack for the DSL line was removed from the router's DSL interface. This loss of network connectivity will result in the IP Phone missing its keepalives with all the phone's configured CallManagers. The 1751 becomes the fallback CallManager and the phone's display will indicate it is in CallManager fall-back mode. This process normally will be completed in approximately two minutes.

At this point, the IP Phone can call an outside phone number by first dialing 9, a second dial tone will be heard, then the area code (if required) and number can be dialed to reach the intended party. Incoming calls will be automatically routed to the extension listed on the **connection plar** [*extension*] command under the voice port. If **connection plar** is not specified, an incoming call will received a dial tone from the Cisco 1751 and the extension can be manually dialed to complete the call. When the router is not functioning as a fallback CallManager, incoming calls will not be completed. It is recommended that no other POTS phones share the analog line connected to the router.

For configurations of SRST based on digital lines and multiple IP Phones, please consult the appropriate configuration and command reference guides.

Design Checklist

This design checklist facilitates pre-implementation planning and the decision process.

Table 4-4 Design Checklist

Design Step	Section References
Identify physical locations for sites to be supported by this design	Organization Specific
Determine IP addressing requirements of branch routers and manual or auto summary scheme	Organization Specific
Decide on location of head-end routers. Will they be in the same rack or across a continent? Will the summary scheme black hole traffic in a failover?	“Head-end Router Locations” section on page 4-24
Determine primary and backup head-end routers. Will branch routers have affinity to the geographically closest head-end?	“Head-end Router Locations” section on page 4-24
Determine number of concurrent voice calls per location. Estimate data bandwidth requirements. Set IP Telephony Call Admission Control requirements appropriately.	“Bandwidth Provisioning for WAN Edge QoS” section on page 4-5
Select appropriate branch site products, based on V ³ PN link speed and other device requirements.	“Branch Office Product Selection” section on page 5-9
Based on the number of remote sites, bandwidth requirements, determine number and location of head-end routers.	“Head-end Product Selection” section on page 5-6
Consider service provider selection process, consult CCO for Cisco Powered Network designated providers.	“Service Provider Recommendations” section on page 4-24
Plan for traffic load sharing requirements for head-end and large branch offices.	“Load Sharing” section on page 4-26
Review existing emergency services plans.	“E911 and 911 Emergency Services” section on page 4-33
Consider SRST requirements for remote locations.	“Survivable Remote Site Telephony” section on page 4-33

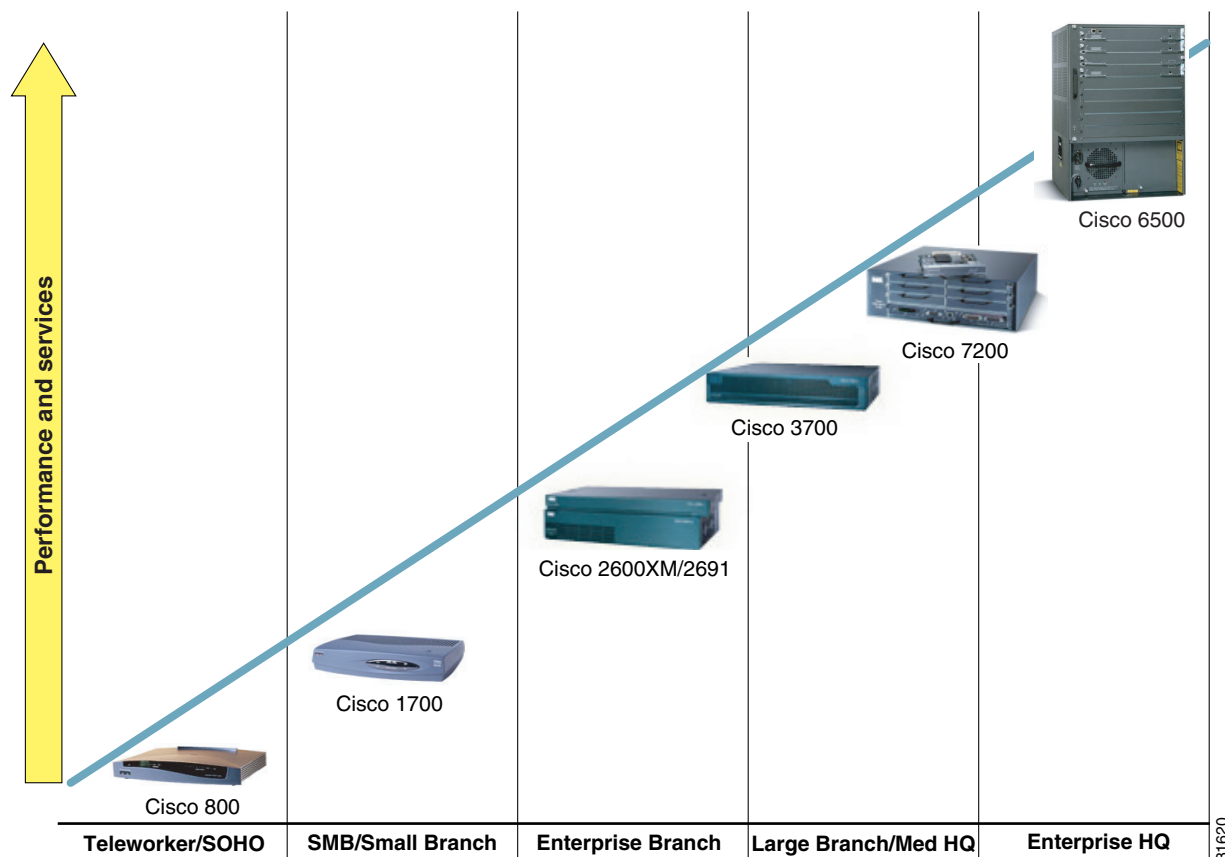


Product Selection

This chapter discusses the V³PN scalability and performance evaluation that was performed and gives product performance results, recommendations, and conclusions that can be used for design parameters when planning and implementing a V³PN deployment.

Cisco offers a whole line of VPN router products which range in application from small business up to large VPN tunnel aggregation points at a large enterprise central site. [Figure 5-1](#) shows the range of products available, almost all of which were evaluated for V³PN performance and deployability.

Figure 5-1 Cisco IOS VPN Router Portfolio



The following topics are addressed within the individual sections of this chapter:

- [Scalability Test Methodology, page 5-2](#)
- [Traffic Profiles, page 5-3](#)
- [Head-end Product Selection, page 5-6](#)
- [Branch Office Product Selection, page 5-9](#)
- [Network Performance/Convergence, page 5-15](#)
- [Software Releases Evaluated, page 5-17](#)

Scalability Test Methodology

The Cisco Enterprise Solutions Engineering VPN performance and scalability lab uses test tools to generate data traffic simulating actual end-to-end applications running on Solaris and Red Hat Linux TCP/IP stacks. The traffic generated incorporates flow control inherit with a TCP implementation. These tools create a network environment that is fairly realistic in terms of how production networks perform.

NetIQ's Chariot test tool is used to generate network traffic. As NetIQ endpoints, SUN NETRA and Penguin Red Hat Linux servers are deployed. The Linux servers generate the simulated voice traffic, the SUN NETRA servers generate the data traffic. Solaris supports path MTU discovery by default.

More information on NetIQ Chariot can be found on the following NetIQ website:

<http://www.netiq.com/products/chr/default.asp>

In addition to the Chariot test tool, portions of the test included implementing two CallManagers, a Survivable Remote Site Telephony (SRST) on a Cisco 2651, and three Cisco 7960 IP phones; one on a campus, one on the SRST Cisco 2651 and one in the core. Using actual phones and voice calls allows for subjective evaluation of the voice quality.

As shown in the diagram in [Appendix A, "Network Diagram Scalability Testbed and Configuration Files"](#), the scalability testbed included 240 branch offices aggregated to two head-end devices. The head-ends consisted of the Cisco 7200 VPN routers (refer to the ["Head-end Product Selection" section on page 5-6](#) for exact models tested). The branch offices consisted of Cisco VPN router products from the Cisco 800, Cisco 1700, Cisco 2600, Cisco 3600, and Cisco 3700 series (refer to the ["Branch Office Product Selection" section on page 5-9](#) for exact models tested).

Branch routers were evaluated at various link speeds ranging from 128 Kbps up to E1. Both head-end and branch router products were evaluated by raising traffic rates up and monitoring key performance parameters until the limitations of each platform were found. The sections that follow give specifics on the exact measurements for each platform. In general key parameters monitored included:

- CPU utilization
- Bi-directional throughput (in bits per second and packets per second)
- End-to-end peak and average latency (for voice traffic)
- End-to-end jitter (for voice traffic)
- Drop rates
- Network resiliency (in the case of head-ends)

All products were evaluated with hardware-accelerated encryption installed. All scalability testing for this Design Guide revision was obtained using IPSec Tunnel Mode, therefore the throughput results might differ in Transport Mode.

In addition to throughput performance testing, failover testing was also conducted. Please refer to the “[Network Performance/Convergence](#)” section on page 5-15 for more information on the failover test scenario.

Traffic Profiles

Cisco Enterprise Solutions Engineering conducted solution testing of this design to validate scalability. Portions of the test plan were designed to simulate worst-case scenarios; intending to find the upper limit or breaking point of the devices under test. Other tests are intended to simulate traffic flows representative of actual production environments.

An initial baseline test was performed in which all traffic was RTP (G.729) streams, without QoS enabled. This was not intended to simulate a specific network, but rather to validate voice quality (latency, jitter and drops would be within acceptable limits) at the upper bounds of the CPU resources of each of the platforms under test.

Then V³PN performance and scalability tests were performed using a converged traffic profile (data and voice) that would be more representative of a real world implementation. In these tests, the traffic profile was as follows:

- UDP—Chariot DNS; script sends 100 bytes in both directions
- UDP—Chariot RTP (VoIPG729) script approximately 33 percent of link capacity (per call units)
- TCP—Chariot HTTPtext script; 100 bytes upstream and 1-to-3K bps downstream
- TCP—Chariot FTPGet and FTPPut script
- TCP—Chariot TN3270
- TCP—Chariot POP3 script; occurs every 1 minute

These different streams were then assigned to realistic QoS traffic classifications as shown in [Table 5-1](#).

Table 5-1 Test Traffic Streams to IP Precedence Mapping

Traffic Stream	IP Precedence
DNS, POP3, FTP	0
RTP (VoIP)	5
HTTP and TN3270	50 percent 2, 50 percent 0

In a production network, EMAIL sent with attachments has characteristics similar to FTP Put or Get, MTU sized packets being sent or received by the user. The POP3 script would represent the text message portion of EMAIL. Also, TN3270 supports screen sizes (lines by columns) of 24x80 (Mod2), 32x80 (Mod3), 43x80 (Mod4) and 27x132 (Mod5). Applications present varying amounts of text on a single screen (transaction), but it is common to see 1-to-2 Kbps downstream, as a 24x80 screen can contain 1920 characters.

Many published test results report performance for IMIX traffic, a traffic pattern developed from packet size samples of Merit’s (Internet) backbone—the “Internet mix” or IMIX. This packet size distribution is typically implemented as 1 packet at 1500 bytes, 4 packets at 512 bytes, and 7 packets at 64 bytes. While this traffic might be representative of the sampled data, it is important for the traffic generation tool to behave as a real TCP or UDP application would on the network. The Chariot test tool used in the

Cisco Enterprise Solutions Engineering lab runs on both Solaris and Linux platforms as end points, and thus uses a real TCP stack, which through adjusting the TCP window size, incorporates a flow control mechanism

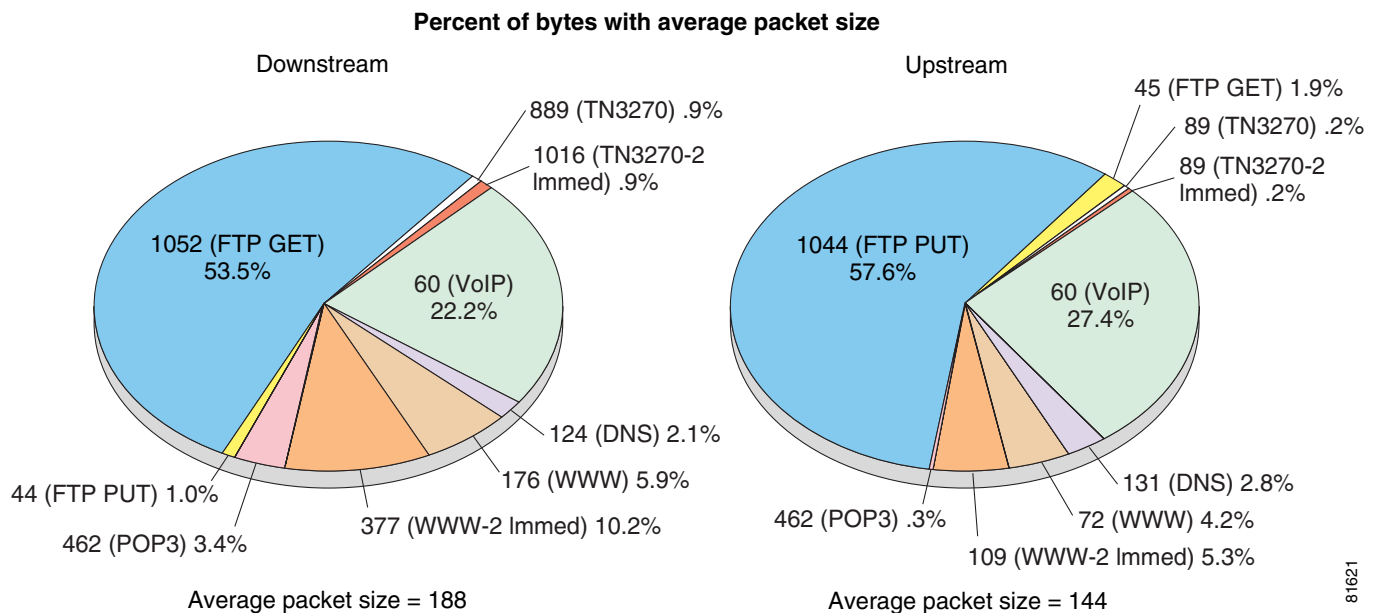
Traffic generation tools that simply generate or replay TCP packets without regard to flow control might facilitate testing, but are not necessarily representative of real applications on a live network.

To monitor and represent the above application mix on a branch under test, a Netflow Protocol-Port-TOS aggregation was configured on a Cisco 2651 router with a WAN link speed of 512 Kbps. Flow exports were captured on a Penguin Linux server for 10 minutes. The raw data was captured and then summarized by application with a Perl script that created a CSV file, which was then graphed. The router under test was configured as follows.

```
ip flow-aggregation cache protocol-port-tos
cache timeout inactive 60
cache timeout active 1
export destination 10.254.0.100 7777
enabled
!
```

Using this captured data, the V³PN performance and scalability traffic profile was verified to be that shown in Figure 5-2. Shown are the different traffic streams, average packet sizes, and percentage of bandwidth consumed for that traffic stream. Note that overhead from IPSec and IP GRE are excluded due to the capture method.

Figure 5-2 V³PN Traffic Mix Sample for 512 Kbps Link Speed



The devices under test have a separate network interface which is used for collecting statistics and configuration. During the tests, the devices are polled via SNMP, are logging to a Syslog server, are configured for NTP, and have telnet sessions active. This management and data collection activity does influence the device's CPU and memory characteristics, but does not represent data traffic on the test network.

Additional Voice Quality Validation

To provide a subjective validation of the Chariot reported values an Agilent Technologies Telegra VQT 2.1 (Voice Quality Tester) was connected to the lab network. The handset cable from a Cisco 7960 IP Phone in a branch office location and the head-end location were connected to the Telegra VQT's handset/audio adapter. An audio file (.wav file) was played by the Telegra VQT through one handset adapter and the resulting voice stream was captured through the handset of the second Cisco 7960 phone.

A test was run on a Cisco 2611 branch router, with hierarchical CBWFQ shaping a T1 interface to a 512 Kbps profile:

```
policy-map 512kb
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  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
!
interface Serial0/0
  bandwidth 512
  ip address 192.168.124.6 255.255.255.252
  load-interval 30
  tx-ring-limit 1
  tx-queue-limit 1
  service-policy output 512kb-shaper
  crypto map static-map
!
```

The Chariot 512 Kbps branch traffic profile was run. This profile would normally include three G.729 RTP streams, but in this case, only two RTP streams were included with the third stream being provided by the Cisco 7960 phone. The test was run for 10 minutes, with the Cisco 7960 dialed using the keypad. Since VAD is by default disabled, the Cisco 7960 generates 50 pps for the entire test/call.

During the test, the blue “i” key (item number 1 in [Figure 5-3](#)) was pressed twice in succession to instruct the Cisco 7960 phone to display the call statistics. Average jitter on the two phones ranged from 8-to-20 msec, maximum jitter on one phone was 45 msec and the second 85 msec, RXDISC on both phones was zero and RXLOST was zero on one phone and five on the second.

Figure 5-3 Cisco 7960 Phone (Illustrating Key Pressed to Display Call Statistics Data)



The Telegra VQT can estimate latency including coder, de-jitter buffer and handset delay—*ear-to-mouth* delay is reported—Chariot does not include or report these values. In the test case the Telegra VQT reported one-way delay to be 122-to-130 msec. If this test configuration would be overlaid to an actual ISP which might include an additional 50-to-60 msec one way delay, it would be practical to assume one way delay less than 200 msec could be achieved.

The test was also run at a 1,024 Kbps data rates with the same Cisco 2611 router, five Chariot G.729 streams plus the Cisco 7960 call for a total of 6 voice streams. The CPU utilization during the test was in the 80 percent range with similar results to the 512 Kbps test.

Head-end Product Selection

This section provides recommendations based on the performance and scalability testing defined above to assist in selecting the appropriate head-end device to meet the VPN aggregation requirements offered by the branch sites.

Failover and Head-end Availability

The primary function of a router is to switch packets. Cisco IOS software has been enhanced to include value add features, for example IPSec, DLSw, TN3270 server. These features are either applications in themselves, like TN3270 server, or support and enhance end user applications—as does IPSec—by encrypting data. Functions such as path determination, which is the job of a routing protocol, are overhead, necessary to support packet switching, but overhead nevertheless.

In Cisco Enterprise Solutions Engineering lab performance testing of the design, a goal of the test was to identify at what rate the head-end routers could switch packets at the highest possible rate, yet still reserve sufficient CPU resources to maintain availability in the event of a network component failure. When redundant network devices fail or are taken out of service for maintenance or upgrades, the path determination process (the routing protocol) generates updates to indicate the network topology change. These updates must be processed by the remaining network components, and the act of processing the updates should not cause other outages in the network.

Consider results observed in one test. Two head-end routers were under test, one actively switching packets and second in reserve with no user data traffic. If the active head-end router was failed, the second router was able to pick up the offered load and maintain all EIGRP neighbor relationships. When the standby router was failed, the active router lost EIGRP neighbor adjacencies as there weren't sufficient reserve CPU cycles to process the topology change. Performance must be balanced with availability requirements.

Performance Under Converged V³PN Traffic Profile

Head-end platforms were configured to aggregate 120 branch offices to the device as primary tunnels and 120 branches as secondary tunnels. The converged (voice and data) V³PN traffic profile discussed in the “Traffic Profiles” section on page 5-3 was then applied and increased while performance of the platform was observed. The upper performance boundary was established by measuring five primary factors:

1. CPU utilization (less than 80 percent)
2. End-to-end RTP packet latency, peak and average (less than 50 msec)
3. End-to-end RTP packet jitter (less than 10 msec)
4. End-to-end RTP bytes lost (less than 0.5 percent)
5. Failover (no loss of peers with failover)

If any factor above was unacceptable, this was considered the breaking point for the product evaluated.

In addition, several other key parameters were also monitored during the evaluation, including:

- Percentage of packets dropped due to IPSec anti-replay and QoS interaction
- Percentage of packets dropped due to the QoS Service Policy
- Portion of packets being process switched
- Occurrences of crypto congestion

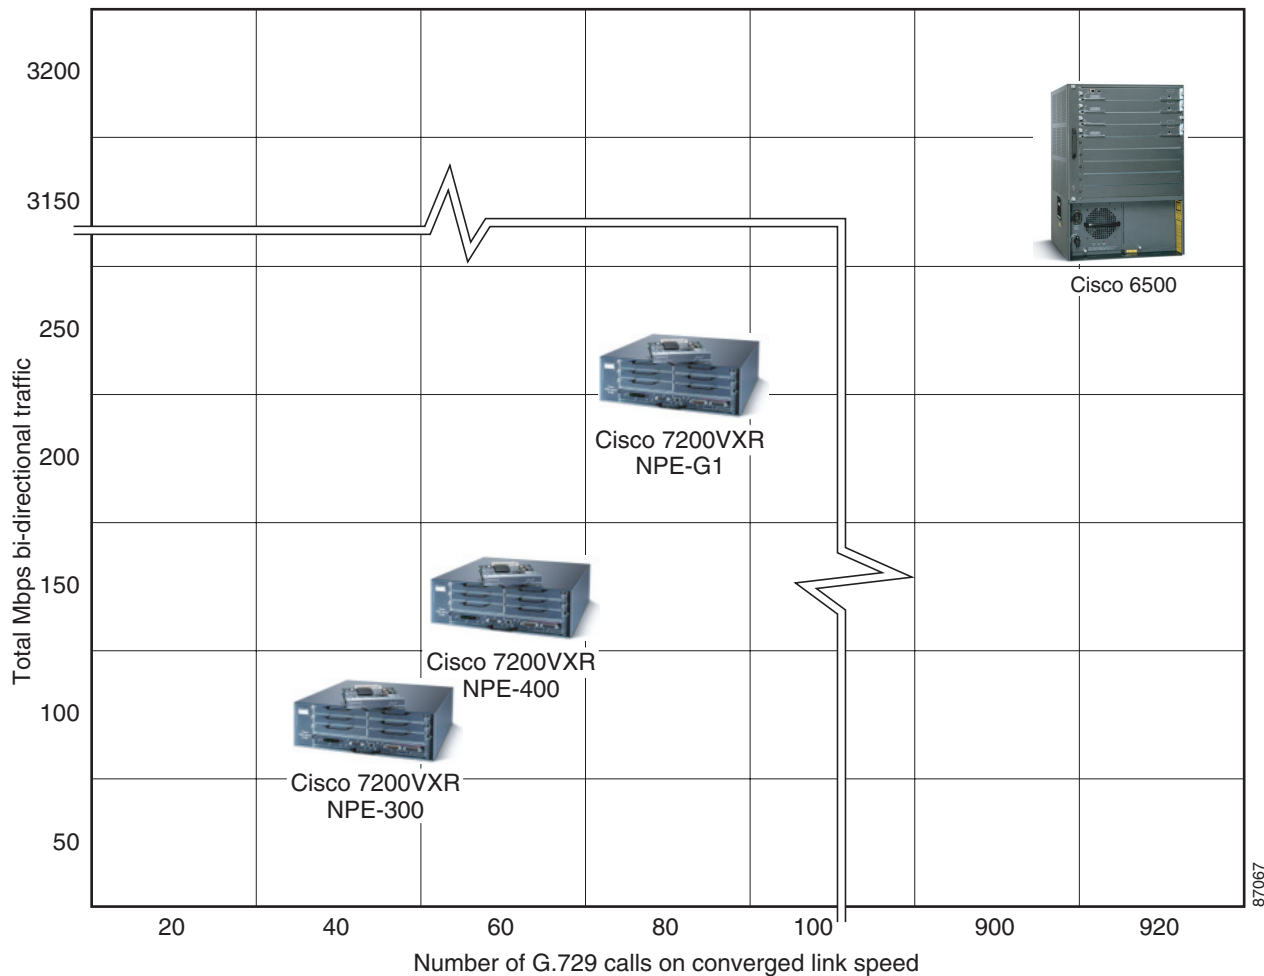
The results of this evaluation are summarized in [Figure 5-4](#).



Note

The scale in [Figure 5-4](#) was modified to accommodate the Cisco 6550 performance results.

Figure 5-4 Head-end Performance—Converged Traffic Case



Impact of QoS on VPN Head-end Performance

The design recommendation for V³PN head-end architecture is to deploy separate devices for WAN aggregation (ISP connection) and VPN tunnel aggregation. WAN aggregation devices would have QoS enabled. The VPN tunnel aggregation devices would typically have Fast Ethernet connectivity to the WAN aggregation devices (no interface congestion).

However, it is possible to have QoS enabled on the same head-end device which is providing VPN tunnel aggregation. This scenario was also evaluated during the scalability and performance testing.

Comparing the relative performance limits at approximately 80 percent CPU shows that having QoS enabled on the VPN head-end device in this evaluation resulted in an approximate 10 percent performance degradation for V³PN throughput. Slightly higher end-to-end latency was also observed, primarily due to the affects of congestion on the output interface since QoS is now operating on the device.

Head-End Scalability and Performance Observations

The following are a summary of the primary observations and findings from the V³PN scalability and performance evaluation of head-end products:

- For both traffic profiles, there is no significant difference in performance between the SA-ISA/ISM and SA-VAM hardware-accelerated encryption cards. The limitation on the Cisco 7200 (NPE-300/NPE-400) is the CPU utilization for forwarding packets.
- As expected, the 7200VXR NPE-G1 was the highest performing Cisco IOS VPN router platform of those evaluated, supporting up to 240 G.729 calls plus approximately 50 Mbps of data in a converged traffic configuration. It should be noted that adding a second VAM to the NPE-G1 platform had very little effect on overall throughput with the converged voice and data traffic profile.
- The Catalyst 6500 with VPN Service Module was also evaluated up to the current 1 Gbps traffic limit of the lab, handling 3,180 G.729 calls and 607 Mbps of simultaneous data traffic. This was in an IPSec-only configuration because GRE currently degrades overall performance of the Catalyst 6500.
- Discounting for the estimated WAN delay in the scalability lab of approximately 4-to-5 msec, end-to-end latency performance was very good: approximately 4-to-5 msec in the voice-only configuration and 15-to-20 msec in the converged traffic configuration.
- Although anticipated to be a major concern, the interaction between IPSec anti-replay and QoS resulted in less than 1 percent of data packets being dropped, primarily due to the positive affect of TCP flow-control. Tuning of the queue-limit parameter in the CBWFQ Service Policy can further reduce drops due to IPSec anti-replay to approximately 0.01 percent of packets.
- The number of VPN tunnels aggregated has significant impact on head-end platform performance. For example, a Cisco 3745 platform running an AIM-II encryption card can process up to 150 G.729 calls over a single tunnel with approximately the same CPU utilization required to process 60 G.729 calls over 60 separate tunnels.
- The impact of enabling QoS on the same VPN head-end device versus on a separate WAN aggregation device resulted in approximately 10 percent performance degradation.
- Monitoring of crypto congestion showed that only under the most extreme cases (for example running the Cisco 7200 above 90 percent CPU utilization) did crypto report congestion. Therefore it is believed with the traffic profile used for evaluation, the LLQ for Crypto feature would not engage, and would not provide additional benefit for head-end platforms.
- However, it is also believed that for organizations with a larger percentage of large packet sizes (such as video), or with applications that might cause *bursts* of large packet sizes, it might be possible to induce congestion on the crypto engine, and in these cases the LLQ for Crypto feature would provide additional insurance that rises in latency did not occur.
- The gating factor for router performance, whether packets are encrypted or clear text, is packets per second, not bytes per second. It requires a similar amount of CPU cycles to encrypt and switch a 64-byte packet as it does a 1400-byte packet. When VoIP is added to the traffic mix, there is a corresponding decrease in average packet size. Consequently, overall throughput in terms of megabits per second is lower with VoIP in the traffic mix because the average packet size decreases.

Branch Office Product Selection

This section provides recommendations based on the performance and scalability testing defined above to assist in selecting the appropriate branch device to meet the VPN requirements for the branch sites.

Due to the large number of branch devices deployed in enterprise networks, the time and expense of installation and upgrades must be balanced with the initial cost and ability to support future network capacity requirements. Many factors besides V³PN performance might need to be considered for branch office router requirements, including:

- What other services will the branch router be performing, such as Cisco IOS Firewall, WAN connection, voice gateway, SRST, or DLSw?
- What is the WAN connectivity, for example point-to-point, DSL, Cable, ISDN, etc.?
- What *future-proofing* factors need to be considered, such as growth potential, new applications on the horizon, etc.?

Product Applicability by Link Speed

This section provides a summary of the ability of different branch platforms to be able to support a given link speed. It can be used as a guideline for deploying the appropriate platform to match up with the network requirements for connection to the service provider.

Table 5-2, Table 5-3, and Table 5-4 summarize the link speeds tested, the maximum number of concurrent G.729 calls for a given link speed, and applicable Cisco VPN router platforms.

Table 5-2 Applicability of Current Products by Link Speed Less Than or Equal to E1

Link Speed (Kbps)	G.729 calls (at 33% of link bandwidth)	831	1751	1760	2611XM	2621XM	2651XM	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)
128	1	x ¹	x	x	x	x	x	x	x	x	x
256	2	x	x	x	x	x	x	x	x	x	x
512	3	x	x	x	x	x	x	x	x	x	x
768	4		x	x	x	x	x	x	x	x	x
1024	6		x	x	x	x	x	x	x	x	x
1280	7		x	x	x	x	x	x	x	x	x
T1 (1536)	9			x		x	x	x	x	x	x
E1 (2048)	12							x	x	x	x

1. x = Supported

Table 5-3 Applicability of Current Products by Link Speed Greater than E1

Link Speed (Mbps)	G.729 calls (at 33% of link bandwidth)	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)	3660 (AIM-II)	2691 (AIM-II)	3725 (AIM-II)	3745 (AIM-II)
3	18	x ¹	x	x	x	x	x	x	x
4	24			x	x	x	x	x	x

Table 5-3 Applicability of Current Products by Link Speed Greater than E1

Link Speed (Mbps)	G.729 calls (at 33% of link bandwidth)	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)	3660 (AIM-II)	2691 (AIM-II)	3725 (AIM-II)	3745 (AIM-II)
5	30					x	x	x	x
10	60					x	x	x	x
15	90								x
20	120								x
25	150								x

1. x = Supported

Table 5-4 Applicability of Legacy Products by Link Speed

Link Speed (Kbps)	G.729 calls (at 33% of link bandwidth)	806	2611	2621	2651	3620	3640
128	1	x ¹	x	x	x	x	x
256	2		x	x	x	x	x
512	3		x	x	x	x	x
768	4			x	x		x
1024	6				x		x
1280	7						
T1 (1536)	9						
E1 (2048)	12						

1. x = Supported

Performance Under Converged V³PN Traffic Profile

Branch platforms were configured with a primary and secondary IPSec/GRE tunnel back to two separate head-ends. The converged (voice and data) V³PN traffic profile discussed in the [“Traffic Profiles” section on page 5-3](#) was then applied and increased while performance of the platform was observed. The upper performance boundary was established by measuring four primary factors:

- CPU utilization (less than 80 percent)
- End-to-end RTP packet latency, peak and average (less than 50 msec)
- End-to-end RTP packet jitter (less than 10 msec)
- End-to-end RTP bytes lost (less than 0.5 percent)

If any factor above was unacceptable, this was considered the “breaking” point for the product evaluated. In addition, several other key parameters were also monitored during the evaluation, including:

- Percentage of packets dropped due to IPSec anti-replay and QoS interaction
- Percentage of packets dropped due to the QoS Service Policy
- Portion of packets being process switched
- Occurrences of crypto congestion

The results of this evaluation are summarized in [Figure 5-5](#) through [Figure 5-8](#).

Figure 5-5 Cisco IOS High-End Branch Router Performance for V³PN

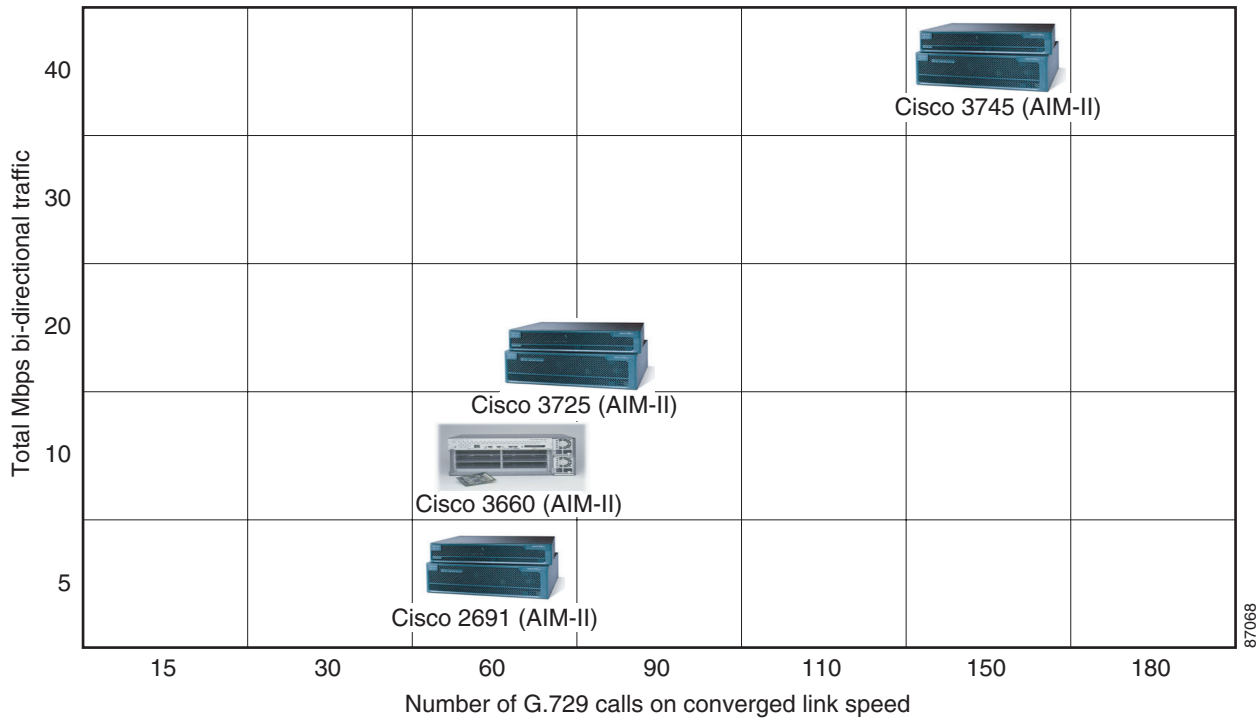
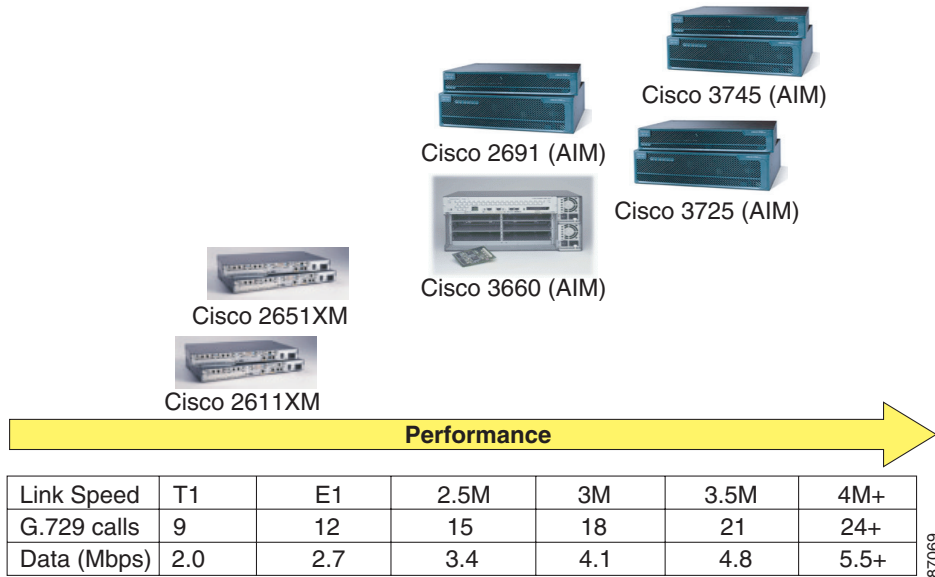
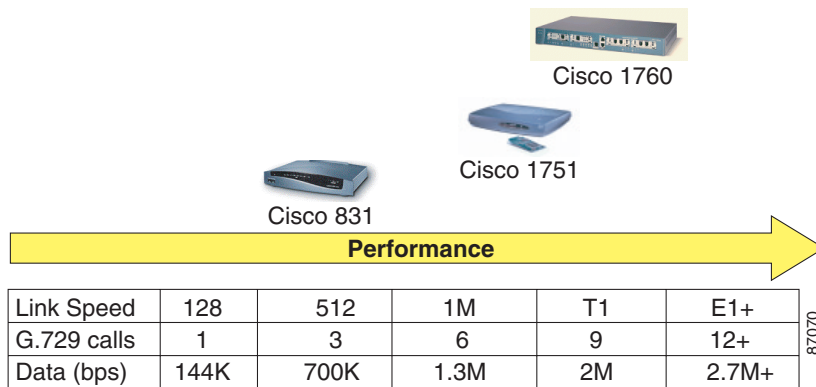


Figure 5-6 V³PN Performance—Current High/Mid-Range Branch Products



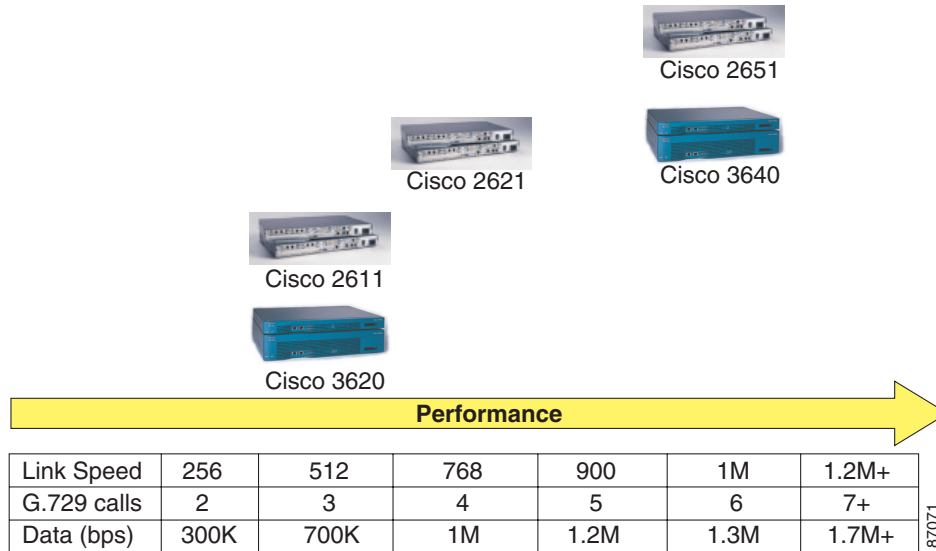
87069

Figure 5-7 Cisco IOS VPN SMB/Small Branch Router—Performance for V³PN



87070

Figure 5-8 Cisco IOS VPN Legacy Branch Router—Performance for V³PN



Branch Scalability and Performance Observations

The following are a summary of the primary observations and findings from the V³PN scalability and performance evaluation of branch products:

- Measured end-to-end latency performance of the solution was very good: approximately 8-to-9 msec in the voice-only configuration and 20-to-25 msec in the Converged traffic configuration (excluding the estimated WAN delay of 4-to-5 ms).
- Although anticipated to be a major concern, the interaction between IPSec anti-replay and QoS resulted in less than 1 percent of data packets being dropped, given the traffic profile discussed earlier, of predominantly TCP-based traffic. Tuning of the queue-limit parameter in the CBWFQ Service Policy can further reduce drops due to IPSec anti-replay to approximately 0.01 percent of packets.
- Monitoring of crypto congestion showed that in many cases crypto did not experience congestion. Therefore it is believed with the traffic profile used for evaluation, the LLQ for Crypto feature would not engage, and would not provide additional benefit for the majority of branch platforms (Cisco 1700, Cisco 2600 and Cisco 3600 VPN routers).
- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers would, because of their much higher CPU capacity, induce congestion on the crypto engine (AIM). Therefore it is believed that the LLQ for Crypto feature would engage on these platforms providing additional insurance that rises in latency did not occur, and that if over-committed, lower priority traffic is dropped, preserving voice traffic. However, with the next generation AIM-II for these platforms, CPU will again be the bottleneck.
- It is also believed that for networks with a larger percentage of large packet sizes (such as video), or with applications that might cause *bursts* of large packet sizes, it might be possible to induce congestion on the crypto engine, and in these cases the LLQ for Crypto feature would provide additional insurance that rises in latency did not occur.
- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers performed well. With current AIM hardware encryption accelerator cards, the Cisco 3660 and Cisco 2691 handled up to 3 Mbps link speeds, while the Cisco 3700s handled up to 4 Mbps link speeds.

87071

- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers with the latest AIM-II hardware encryption accelerator cards performed exceptionally well. The Cisco 3660, Cisco 2691, and Cisco 3725 handled up to 60 simultaneous calls on 10 Mbps link speeds, while the 3745 handled up to 150 simultaneous calls at a 25 Mbps link speed. Keep in mind that testing was performed on these platforms in a branch VPN router scenario, meaning two IPSec/GRE tunnels (primary and secondary) established to the head-end VPN routers. If deployed as a hub/head-end tunnel aggregation device, performance would be expected to decrease as the number of IPSec/GRE tunnels is increased.
- The Cisco 806 VPN router has some limitations if deployed as a V³PN device. Since the Cisco 806 does not support hardware-accelerated encryption, latency performance can quite easily be disrupted by data traffic or issuing CLI commands and cause voice latency spikes in a range of 500 msec to 1 second. Therefore, deploying this platform is not recommended. The Cisco 831 should be deployed instead.
- The Cisco 1751/1760 VPN routers performed quite well, handling up to a T1 link speed of converged traffic. However, the Cisco IOS software revision required (containing QoS Pre-Classify) is a special branch, 12.2(4)YB. Also, process switching was occurring on the majority of packets on the 1700.
- The Cisco 1721 VPN router was not evaluated; however, since it is similar in CPU and VPN hardware-acceleration, V³PN performance should be analogous to Cisco 1751 and Cisco 1760 performance.
- The legacy Cisco 2600/3600 VPN routers were found to have limitations when subscribed to higher link speeds (such as T1) due to the nature of their packet switching characteristics. Therefore the recommendation is to follow the guidance given in the [“Product Applicability by Link Speed” section on page 5-10](#).
- The Cisco 2600XM VPN routers were evaluated and do not appear to have the limitations identified on older Cisco 2600 VPN router platforms. Cisco 2600XM platforms handled up to T1 link speeds of converged V³PN traffic.
- The scalability evaluation was performed with both Frame Relay and HDLC as the L2 encapsulation. Performance was fairly comparable with HDLC providing slightly higher link utilization.

Network Performance/Convergence

Each organization might have different convergence time requirements. The design principles in this guide were used to perform a scalability test with 240 branch offices aggregated to two Cisco 7200 NPE-400 head-end devices.

Two aggregation configurations were evaluated:

- Active/Standby—One head-end device was loaded with all 240 primary IPSec/GRE tunnels (170 with active traffic streams), while the second is in a *standby* mode configured with all 240 secondary tunnels.
- Active/Active—Each of the two head-end devices is configured with 120 primary (85 with active traffic streams) and 120 secondary IPSec/GRE tunnels, fairly equal loading on both head-ends.

The test was performed by powering off one of the head-end devices to simulate a complete failure. In this test, the network fully converged after approximately 20-to-23 seconds. The starting and failover traffic/tunnel aggregation conditions are shown below in [Table 5-5](#).

Table 5-5 Head-end Failover Scenario

	Head-end 1	Head-end 2
Active-Standby Active Fails		
Starting Condition	52.0 Mbps total traffic 170 calls 240 primary tunnels 80 percent CPU	120 Kbps total traffic 0 calls 240 secondary tunnels 1 percent CPU
During Simulated Failure	Simulated Failure	61.2 Mbps total traffic 170 calls 240 primary tunnels 79 percent CPU
Active-Standby Standby Fails		
Starting Condition	51.7 Mbps total traffic 170 calls 240 primary tunnels 76 percent CPU	120 Kbps total traffic 0 calls 240 secondary tunnels 1 percent CPU
During Simulated Failure	60.3 Mbps total traffic 170 calls 240 primary tunnels 83 percent CPU	Simulated Failure
Active-Active Active Fails		
Starting Condition	25.9 Mbps total traffic 85 calls 120 primary tunnels 120 secondary tunnels 46 percent CPU	25.8 Mbps total traffic 85 calls 120 primary tunnels 120 secondary tunnels 46 percent CPU
During Simulated Failure	Simulated Failure	50.2 Mbps total traffic 170 calls 240 primary tunnels 80 percent CPU

**Note**

The traffic and voice call loads were intentionally raised near or slightly above recommended operating limits to characterize worst-case failover conditions. It is not recommended to operate platforms at the extremes shown in this table.

While all VoIP streams experienced packet loss during the failure, after routing convergence all simulated call streams continued to function within acceptable end-to-end latency, jitter and packet loss limits.

In all scenarios, the failed head-end device was then powered back on, resulting in the network re-converging in less than two seconds. The IPSec tunnels re-established a few at a time as their corresponding SAs were renegotiated. The last IPSec tunnels re-established connectivity after 3-to-4 minutes.

Software Releases Evaluated

The software releases shown in [Table 5-6](#) were used in the V³PN scalability and performance evaluation.

Table 5-6 Software Releases Evaluated

Cisco Product Family	SW Release	Notes/Other Information
Cisco 7200VXR VPN Routers	Cisco IOS software 12.1(9)E	3DES IPsec support C7200-IK2S-M
Cisco 3700 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C3725-IK9O3S-M
Cisco 3600 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C3640-IK9O3S-M
Cisco 2600 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C2600-IK9O3S-M
Cisco 1700 Series VPN Routers	Cisco IOS software 12.2(4)YB	3DES IPsec support C1700-K9O3SY7-M
Cisco 800 Series VPN Routers	Cisco IOS software 12.2(8)YN	3DES IPsec support C831-K9O3SY6-M
Cisco 7500 WAN Aggregation Routers	Cisco IOS software 12.2(4)XV4	RSP-JSV-M
Cisco 6x00 Catalyst Switches	CatOS 5.5(3)	
Cisco 2948G Catalyst Switches	CatOS 4.5(9)	



Note

Several Cisco IOS software images exist, each configured with various levels of encryption technology. There are certain restrictions and laws governing the use and export of encryption technology.

Before selecting Cisco IOS software, perform the appropriate research on www.cisco.com and consult the appropriate support channels. It is important understand issues inherent to specific levels of Cisco IOS software code that might affect other features configured on the network's routers.



Implementation and Configuration

This chapter provides step-by-step examples of how to configure the V³PN environment tested by Cisco Enterprise Solutions Engineering and is separated into the following principal sections:

- [Routing Protocol, Switching Path and IP GRE Considerations, page 6-1](#)
- [QoS Configuration, page 6-5](#)
- [WAN Aggregation Router Configuration, page 6-9](#)
- [IKE and IPsec Configuration, page 6-19](#)

The chapter ends with the “[Implementation and Configuration Checklist](#)” section on [page 6-24](#).

The recommended approach would be to read through each section before implementing V³PN. If there are topics or concepts which are unclear, consult the associated design guides or documentation at Cisco’s informational website at <http://www.cisco.com> for more details. Then approach the first configuration with the Implementation and Configuration checklist and refer back to the specific section for examples.

Routing Protocol, Switching Path and IP GRE Considerations

This SRND assumes that the IP addressing scheme lends itself to summarization from the branch to the core, and that the core sends a default or summary route to the branch. Design considerations specific to this implementation are detailed in the following sections:

- [Configure Switching Path, page 6-1](#)
- [Configure IP GRE Tunnels, page 6-2](#)
- [EIGRP Summarization and Network Addressing, page 6-2](#)
- [EIGRP hold-time, page 6-3](#)
- [IP GRE Tunnel Delay, page 6-3](#)

Configure Switching Path

In the Cisco Enterprise Solutions Engineering lab test, CEF was enabled on all routers under test. To enable CEF switching on the branch routers, verify or configure the `ip cef` command.

Configure IP GRE Tunnels

When configuring IP GRE tunnels, the most common issue is routing. The IP addresses selected for the tunnel end-points—the tunnel source and destination—must be reachable for the tunnel to come up. Additionally, when a routing protocol is configured on the tunnel interface, it must not advertise the network(s) represented by the tunnel end-points through the tunnel itself or the tunnel interface will be disabled and the following message is logged:

```
%TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing
```

Please see www.cisco.com for more information on configuring tunnel interfaces. Specifically,

http://www.cisco.com/warp/public/105/gre_flap.html

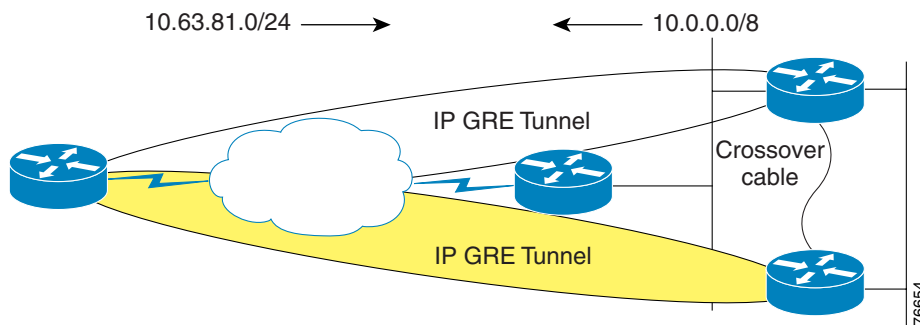
contains information tunnel interfaces and routing. There is also a complete illustration in [Appendix A, “Network Diagram Scalability Testbed and Configuration Files”](#) of this publication.

To avoid problems when configuring IPSec, configure and verify both tunnel interfaces are up to the head-end IPSec/GRE head-end routers.

EIGRP Summarization and Network Addressing

In this phase of Cisco Enterprise Solutions Engineering lab testing, EIGRP Stub support was not scale tested. The EIGRP Stub feature increases scalability and it will be incorporated into future full-scale testing. Optimized addressing and summarization are implemented in the design. With or without implementing EIGRP Stub support, proper summarization enhances network stability and performance. Just as IPSec increases the bandwidth requirements of voice packets, it also adds to the bandwidth required for routing protocol updates. Decreasing the number of routing updates that must be sent is under control of the network manager, and summarization is the tool used to implement that bandwidth savings. [Figure 6-1](#) illustrates one method of IP addressing and EIGRP summarization.

Figure 6-1 IP Addressing and EIGRP Summarization



Branch routers advertise a manually summarized route on a 24 bit boundary via the tunnel interfaces to the core:

```
!
interface Tunnel0
 ip summary-address eigrp 1 10.63.81.0 255.255.255.0 5
!
interface Tunnel1
 ip summary-address eigrp 1 10.63.81.0 255.255.255.0 5
!
```

Core routers advertise a manually summarized route on an 8-bit boundary via the tunnel interfaces to the branch routers:

```
!
interface Tunnel240
 ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
!
```

The typical branch EIGRP configuration is shown below:

```
!
router eigrp 1
 passive-interface Serial0/0.1
 passive-interface Ethernet0/1
 network 10.0.0.0
 no auto-summary
 eigrp log-neighbor-changes
!
```

EIGRP hold-time

On 64 Kbps links in the Cisco Enterprise Solutions Engineering lab tests the hold-time on the tunnel interfaces was increased from the default of 15 seconds to 25 seconds to maintain EIGRP neighbor relationships in that event that three consecutive EIGRP hello packets were dropped.

```
!
interface Tunnel0
 ip hold-time eigrp 1 25
!
```

The hello interval remained at the default of 5 seconds. While EIGRP hello packets have their PAK_PRIORITY bit set to indicate relative importance on the originating router, once EIGRP hello packets are encapsulated in IP GRE and IPsec headers, the only indication of significance to intermediate routers (ISP routers) is the IP Precedence of 6. This underscores the importance of using a QoS aware service provider.

IP GRE Tunnel Delay

The delay value was increased on the backup tunnel interface (Tunnel 1 on the branch routers) to influence path selection with EIGRP. The default delay value for a GRE tunnel interface is 500000 usec, to make one tunnel the backup interface the delay value was increase to 600000 usec. All traffic traverses the primary tunnel unless the head end device is unavailable. In the Cisco Enterprise Solutions Engineering lab test, there is only one physical interface and the tunnels are sourced off the physical interface. If there were two physical interfaces per branch, it would be preferable to source off loopback interfaces so both logical tunnels remain up in the event of a branch serial interface failure.

The **show interface** commands displays delay in microsecond units. The **delay** interface command specifies the delay metric, in 10 microsecond units. EIGRP calculates its metric from the minimum bandwidth in Kbps for all links in the path, and the cumulative delay in microseconds for all links in the path.

```

!
vpn13-3640-2#show interface tunnel 0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  Description: Tunnel0
  Internet address is 10.63.81.194/30
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
    reliability 255/255, txload 255/255, rxload 255/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 192.168.217.2, destination 192.168.252.1
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Tunnel TTL 255
  Checksumming of packets disabled, fast tunneling enabled

!
!
interface Tunnel0
  description Tunnel0
  ip address 10.63.81.194 255.255.255.252
  ip hold-time eigrp 1 25
  ip summary-address eigrp 1 10.63.81.0 255.255.255.0 5
  load-interval 30
  tunnel source 192.168.217.2
  tunnel destination 192.168.252.1
  crypto map static-map
!
vpn13-3640-2#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  Description: Tunnel1
  Internet address is 10.63.81.198/30
  MTU 1514 bytes, BW 9 Kbit, DLY 600000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 192.168.217.2, destination 192.168.251.1
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Tunnel TTL 255
  Checksumming of packets disabled, fast tunneling enabled

!
!
interface Tunnel1
  description Tunnel1
  ip address 10.63.81.198 255.255.255.252
  ip hold-time eigrp 1 25
  ip summary-address eigrp 1 10.63.81.0 255.255.255.0 5
  delay 60000
  tunnel source 192.168.217.2
  tunnel destination 192.168.251.1
  crypto map static-map
!
!
interface Serial0/0.1 point-to-point
  description Serial0/0.1
  ip address 192.168.217.2 255.255.255.252
!

```


QoS Configuration

Quality of Service (QoS) implementation topics covered in this chapter are:

- [Campus QoS—Mapping ToS to CoS, page 6-5](#)
- [QoS Trust Boundary, page 6-6](#)
- [Configure QoS Class Map, page 6-6](#)
- [QoS Policy Map Configuration, page 6-7](#)

Campus QoS—Mapping ToS to CoS

Use separate VLANs for voice and data when there is an option to segment the IP address space at the branch office. If the switch in use at the branch supports only Layer-2 services, no Layer-3, and supports 802.1Q trunking, then the branch WAN router should be configured to set the User Priority bits in the 802.1p portion of the 802.1Q header. The **set cos** is only supported with IEEE 802.1Q/ISL interfaces.

The Cisco 800 series, 1720 and 1750 do not support 802.1Q, however the Cisco 2600 and 3600 series, as well as the Cisco 1710, 1751, 1760 and 1721 do support 802.1Q. The IP Plus feature set is needed to support 802.1Q trunking.

The following example should be used to supplement the configuration files from the lab test results.

```
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
class-map match-all user-mission-critical
  match ip precedence 2
!
policy-map output-L3-to-L2
  class voice
    set cos 5
  class call-setup
    set cos 3
  class user-mission-critical
    set cos 2
!
interface FastEthernet0/1.201
  encapsulation dot1Q 201
  ip address 10.250.0.1 255.255.255.0
  service-policy output output-L3-to-L2
!
```

In this example, a **user-mission-critical** class is used to specify only the IP Precedence 2 mission critical traffic that is bound for the end-user's workstation via the branch router's Fast Ethernet interface. The combination of IP Precedence 2 and 6 traffic into the **mission-critical** was intended to provide a class for both end-user mission critical traffic as well as EIGRP hello/updates and other management traffic like SNMP, Telnet, NTP, etc. which are set to IP Precedence 6.

For additional information, refer to the *Enterprise QoS Design Guidelines* at the following location:

http://www.cisco.com/application/pdf/en/us/guest/netsol/ns17/c649/ccmigration_09186a00800d67ed.pdf

QoS Trust Boundary

In this design guide, it is assumed that the network manager is setting the IP Precedence/DSCP values appropriately so they match the service policy applied to the output interface. This can be done at the application level, or the Layer-2 switch at the remote and head-end. For the purpose of simplicity in our lab testing, it is assumed IP Precedence is marked by the application or Layer-2 switch. Cisco IP Phones set the IP precedence of the voice traffic to 5 (DSCP *EF*) and override the IP Precedence of data traffic from the switch port in the phone to 0.

Not all devices are attached to an IP Phone. An example would be a dedicated DLSw router at the branch or DLSw configured on the branch WAN router. It is extremely important to audit the applications in use against the policy-map implemented.

DLSw peers generate TCP traffic to/from port 2065 and by default DLSw sets IP Precedence to 5. In the DLSw configuration change this by using `dlsw tos map`, if priority peers are not configured, high is the only operative option. For example:

```
!
dlsw local-peer peer-id 10.251.0.1
dlsw remote-peer 0 tcp 10.254.0.45
dlsw tos map high 2 medium 0 normal 0 low 0
!
```

Now DLSw is set to use IP Precedence of '2' and its traffic matches on the **mission-critical** entry, rather than the 'voice' class, which it would if left to the default.

The [“Using NetFlow to Verify ToS Values” section on page 7-6](#) presents an example of using NetFlow to verify the ToS byte values used by applications on the network.

Configure QoS Class Map

In order to configure a QoS Policy map, a class map must first be configured. The purpose of the class map is to define the packets associated with the named class of traffic. Just as each enterprise is different in the type and nature of its applications, so will the class map definition. The following configuration sample was used during lab testing.

```
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
!
```

Expect to *tune* the **class-map** configuration during implementation. Business critical applications have a way of being overlooked until they don't work properly. In the above example, a **mission-critical** class was created and it includes both user mission critical applications (IP Precedence 2) and *Internetwork Control* or IP Precedence 6 traffic. Cisco telnet, BGP, EIGRP, OSPF, NTP, SNMP all use IP Precedence 6 and are included in this class. During testing, EIGRP hellos were being dropped—they defaulted to the **class-default** so IP Precedence 6 was included in the **mission-critical** class. Another approach would have been to create a distinct class for this traffic.

For example:

```
!  
class-map match-all call-setup  
  match ip precedence 3  
class-map match-all mission-critical  
  match ip precedence 2  
class-map match-all internetwork-control  
description Routing hellos/updates, cisco telnet, SNMP, NTP  
  match ip precedence 6  
class-map match-all voice  
  match ip precedence 5  
!
```

However, this configuration was not scale tested in the lab. The amount of bandwidth to allocate to the **internetwork-control** class can vary depending on parameters such as the degree of summarization of routing protocol advertisements, the volume of SNMP, or telnet traffic. A suggested starting value is allocating five percent of the bandwidth to **internetwork-control** and increasing that as needed depending on the implementation.

The QoS Class Map must be configured on each organization's branch router. In addition, if a Layer-3 service provider is being used, the QoS Class Map must also be configured on the service provider's edge router.

For the central site, the QoS Class Map can be configured on either the enterprise head-end WAN aggregation routers (in the case of separate WAN aggregation and VPN tunnel aggregation) or on the VPN head-end devices (in the case of no separate WAN aggregation device). No QoS need be configured on the VPN head-end routers if they are Fast Ethernet in and out (i.e. separate WAN aggregation and VPN head-ends).

QoS Policy Map Configuration

In the design section of this guide the assumption was made that there are four types or categories of traffic on the network:

- Voice Bearer (VoIP RTP packets)
- Voice Control (Call Control Signaling)
- Mission Critical (End-user and Internetwork Control traffic)
- All other Data

This traffic is identified by the value of the ToS byte (IP Precedence or optionally DSCP), the routers are not matching on port, protocol or IP addresses, although this is an alternative in networks which are not end to end QoS enabled by the application hosts.

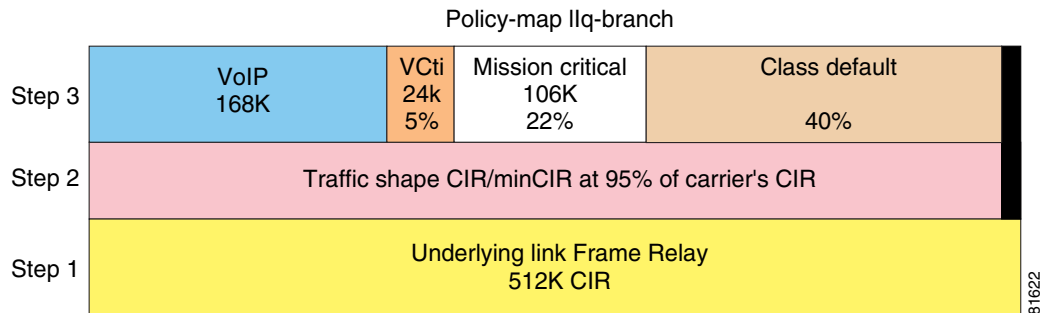
Configuration Example—512 Kbps Branch

For illustration purposes, assume a 512-Kbps CIR Frame Relay PVC is provisioned at this site. Up to three concurrent voice calls are active—three at 56 Kbps per call or 168 Kbps for voice bearer. Voice Control is allocated 5 percent of the underlying bandwidth and Mission Critical 22 percent.

To avoid an increase in latency as the network traffic approaches the carrier's CIR, traffic shape to 95 percent of the carrier's CIR.

Using a 512 Kbps link as an example, review [Figure 6-2](#) as an illustration of how the policy map overlays on the provisioned link.

Figure 6-2 Example Bandwidth Provisioning for 512 Kbps



The link in the example presented in Figure 6-2 is provisioned as follows:

1. The carrier provisions the link with a Frame Relay CIR of 512 Kbps.
2. The Frame Relay **map-class** configuration will traffic shape to 95 percent of the carrier's CIR.
3. For serial interfaces with HDLC encapsulation, the priority class's Kbps is added to the sum of the percent classes. If the resulting value is within 75 percent (default value for **max-reserved-bandwidth**) of the interface's bandwidth, the allocation is accepted. The **max-reserved-bandwidth** command is not supported on Frame Relay PVCs; however it would be recommended to provision the priority and bandwidth classes total allocated bandwidth within 75 percent of the underlying link bandwidth.

The following is a sample configuration implementing the traffic categories for a 512 Kbps branch:

```

!
hostname vpn9-2600-1
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
!
!
policy-map llq-branch
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority 168
  class class-default
    fair-queue
!
end

```

For link speeds other than 512 Kbps, modify the kilobits per second value of the voice class's priority parameter. If separating **internetnetwork-control** into a separate class, allocate 5 percent initially and monitor for drops in this class and increase as needed. For example:

```
!
class-map match-all mission-critical
  match ip precedence 2
class-map match-all internetnetwork-control
  match ip precedence 6
!
policy-map llq-branch
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class internetnetwork-control
    bandwidth percent 5
  class voice
    priority 168
  class class-default
    fair-queue
!
```

This configuration must be applied to the branch router, the WAN aggregation routers and the corresponding service provider's routers if using a Layer-3 service provider.

Please refer to the “[Anti-Replay Considerations](#)” section on page 4-16 for an illustration of modifications that can be made to the **queue-limit** parameter within bandwidth classes to reduce anti-replay drops.

WAN Implementation Considerations

In deploying the V³PN implementation presented in this SRND, the following key discussions are provided:

- [WAN Aggregation Router Configuration](#), page 6-9
- [Frame Relay Traffic Shaping and FRF.12 \(LFI\)](#), page 6-11
- [Attach Service Policy to Frame Relay Map Class](#), page 6-14
- [Apply Traffic Shaping to the Output Interface](#), page 6-15
- [Applying Service Policy to HDLC Encapsulated T1 Interfaces](#), page 6-16
- [Combined WAN and IPSec/IP GRE Router Configuration—Cisco 7200 HDLC/HSSI](#), page 6-17

WAN Aggregation Router Configuration

In this publication, the preferred implementation separates the IPSec/IP GRE head-end routers from the WAN aggregation routers. The following configuration example is for a 512 Kbps link from the WAN aggregation router., in this case a Cisco 75xx series with VIP4-80 and Channelized T3 interfaces. This Distributed Traffic Shaping configuration offloads the traffic shaping function from the Route Switch Processor (RSP) to the VIP. Distributed CEF (dCEF) was also configured.

The following configuration is an illustration of one link/subinterface. There must be hierarchical policy-maps for *each* different link speed represented. Each branch in this configuration would have its own time slots and subinterfaces.

```

vpn2-7500-2

1 GEIP controller (1 GigabitEthernet).
3 VIP4-80 RM7000 controllers (4 Serial)(5 Channelized T3).
1 Gigabit Ethernet/IEEE 802.3 interface(s)
124 Serial network interface(s)
5 Channelized T3 port(s)
!
ip cef distributed
!
controller T3 2/0/0
  clock source line
  cablelength 50
  t1 1 channel-group 1 timeslots 1-4
  t1 2 channel-group 1 timeslots 1-4
  t1 3 channel-group 1 timeslots 1-4
  t1 4 channel-group 1 timeslots 1-4
  t1 5 channel-group 1 timeslots 1-4
  t1 6 channel-group 1 timeslots 1-4
  t1 7 channel-group 1 timeslots 1-16
  t1 8 channel-group 1 timeslots 1-2
  t1 9 channel-group 1 timeslots 1-4
  t1 10 channel-group 1 timeslots 1-8
! Eight timeslots at 64K each = 512Kbps
[...]
t1 28 channel-group 1 timeslots 1-2
!
policy-map 512kb
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority 168
  class class-default
    fair-queue
policy-map 512kb-shaper
  class class-default
    shape average 480000 1920 0
    service-policy 512kb

!
interface Serial2/0/0/10:1
  description vpn13-2600-4
  no ip address
  encapsulation frame-relay
  no fair-queue
!
interface Serial2/0/0/10:1.102 point-to-point
  description vpn13-2600-4
  ip address 192.168.214.1 255.255.255.252
  frame-relay interface-dlci 102
  class 512kb
!
map-class frame-relay 512kb
  no frame-relay adaptive-shaping
  service-policy output 512kb-shaper
  frame-relay fragment 640
!
end

```

Frame Relay Traffic Shaping and FRF.12 (LFI)

Frame Relay traffic shaping is configured to rate limit the output packets to the carrier's provisioned CIR as opposed to the clock rate of the output interface. In the case of a branch router, the serial interface clock rate (port speed) might be at a T1 line rate, but the CIR would be provisioned at 768 Kbps. At the head-end WAN aggregation router the difference between output interface might be even greater, an individual site's 128 Kbps CIR might connect to the carrier over a HSSI or Channelized T3 interface.

To eliminate the serialization delay for voice packets on low speed links, Link Fragmentation and Interleaving (LFI) should be configured on interfaces less than 1024 Kbps. On Frame Relay encapsulated interfaces, LFI is implemented by FRF.12.

The key Frame Relay considerations are:

- FRF.12 is configured on a per-PVC basis
- Frame Relay traffic shaping must be enabled on the interface for fragmentation to work
- This design implements fragmentation byte counts consistent with 10 msec of delay
- When FRTS and FRF.12 are enabled a dual interface FIFO queue is created, each PVC's priority queue and LMI packets go to the high queue, all other classes into the normal queue.

The following configuration example should be used as a template for the branch router configuration. In this configuration, each branch has **one** Frame Relay PVC to the service provider. To increase availability (and cost), two PVCs could be configured—one to each head-end WAN aggregation router. Another alternative would be to provision two serial interfaces, one each to separate Frame-relay providers terminated on separate head-end WAN aggregation routers. A sample branch router with a CIR of 512 Kbps follows:

```
!
hostname vpn13-1700-4
!
interface Serial1/0
  description Serial1/0
  bandwidth 512
  no ip address
  encapsulation frame-relay
  logging event subif-link-status
  logging event dlci-status-change
  load-interval 30
  frame-relay traffic-shaping
  frame-relay lmi-type cisco
!
interface Serial1/0.1 point-to-point
  description Serial1/0.1
  bandwidth 512
  ip address 192.168.224.2 255.255.255.252
  frame-relay interface-dlci 101
    class ts-branch
  crypto map static-map
!
map-class frame-relay ts-branch
  frame-relay cir 486400
  frame-relay bc 4864
  frame-relay be 0
  frame-relay mincir 486400
  no frame-relay adaptive-shaping
  service-policy output llq-branch
  frame-relay fragment 640
!
end
```

Table 6-1 summarizes the different parameters available for Frame Relay Traffic Shaping.

Table 6-1 Frame Relay Traffic Shaping Parameters

Parameter	Explanation
no frame-relay adaptive-shaping	FRTS will not decrease its sending rate based on receipt of BECN or ForeSight backward congestion notification messages. However, the target rate is MINCIR for adaptive shaping, and MINCIR = CIR in this design.
frame-relay cir	The value is carrier's CIR * 0.95 (rounded down), so the carrier's switch will not see the router sending at or above CIR.
frame-relay bc	For VIP based routers router's CIR * 0.004, for low end systems (branch routers) router's CIR * 0.01
frame-relay mincir	Same as router's CIR. The default is CIR/2.
frame-relay be	Default value of 0

Table 6-2 illustrates parameter values that can be used in the above configuration for the link speeds in the Cisco Enterprise Solutions Engineering lab test.

Table 6-2 Frame Relay Traffic Shaping Parameters

Line Rate (Kbps)	Cisco IOS TS CIR/minCIR (BW*.95)	Cisco IOS TS Bc (CIR*.01)	LFI bytes	VIP CIR/MinCIR (BW*.95) down to 8000 multiple	VIP Bc (CIR*.004)
64	60800	608	80	56000	224
128	121600	1216	160	120000	480
256	243200	2432	320	240000	960
512	486400	4864	640	480000	1920
768	729600	7296	1000	728000	2912
1024	972800	9728	N/A	968000	3872
1280	1216000	12160	N/A	1216000	4864
1536	1459200	14592	N/A	1456000	5824
2048	1945600	19456	N/A	1944000	7776
3072	2918400	29184	N/A	2912000	11648
6144	5836800	58368	N/A	5832000	23328

Frame Relay adaptive shaping is targeted for a configuration where the CIR value equates to the port speed and the MINCIR value is the carrier's CIR value. This configuration allows the network to burst to port speed when no congestion exists in the carrier's network but to traffic shape to CIR during periods of congestion. While this might be advantageous for data only environments, it is not recommended for converged voice and data networks, so it is disabled.

The router's CIR value is shown as 95 percent of the carrier's CIR value. This is to eliminate the possibility of sending data at or above CIR from the switches perspective. The 95 percent is a conservative approach, to prevent over-subscription if the router and frame switch account for Layer-2 overhead differently.

Frame Relay Traffic Shaping involves a concept of “metered bursting”, where during an interval of time some number of bits can be sent (or burst) into the Frame Relay carrier’s network. The numbers of bits are specified as the committed burst (Bc) and this number of bits is divided by the CIR, or average rate, to derive an interval of time. $\text{Interval} = \text{Bc}/\text{CIR}$. The Cisco default Frame Relay shaping parameters are Bc is 1/8 of CIR. This default value is 125ms. This value is optimized for data traffic, but introduces delay for voice packets. Simply stated, it is possible the committed burst number of bits transmitted will be exhausted in the first 5 msec of the interval, and thus the algorithm will wait 120 msec before transmitting a subsequent burst.

To optimize the Frame Relay Traffic Shaping parameters—average rate, Bc, and excess burst (Be) for voice—the interval size is reduced. A smaller interval size equates to more intervals per second. An interval size optimal for voice would be in the 10-to-20 msec range or a Bc value of 1-to-2 percent of CIR.

There is, however, a negative effect of this optimization for voice. By decreasing the Bc value, Frame Relay Traffic Shaping becomes engaged or *active* more aggressively. This in turn provides congestion feedback to the CBWFQ service policy and it might drop or delay packets before the average rate approaches the CIR. This in turn can delay data packets and trigger anti-replay drops. This symptom was exhibited when a show interface, using a 30-second load interval, reports 50-to-60 percent utilization.

The Frame Relay Bc value reduces the interval from the default of 125 msec to 10 msec. Normally Bc is 1/8 of the CIR value which equates to 125 msec. A Bc value of 1 percent the CIR configures an interval of 10 msec for all line rates. For example:

```
vpn18-2600-2#show traffic-shape s0/0.100
```

Interface	Se0/0.100	Access Target	Byte Limit	Sustain bits/int	Excess bits/int	Interval (ms)	Increment (bytes)	Adapt Active
VC	List	Rate	Limit	bits/int	bits/int	(ms)	(bytes)	Active
100		486400	608	4864	0	10	608	-

The nature of traffic shaping is to delay, or buffer packets, so the sending rate equates to the credit of bytes per interval. If the Interval is left at the default value of 125 msec, bursts of data traffic that exceed the credit cause the traffic shaping algorithm to wait for that interval to expire before attempting to transmit subsequent packets (which could be voice packets). Decreasing Bc to 10 msec is an accommodation to maintain voice packets within their delay budget.

The definition of *MINCIR* is the minimum amount of data to be sent during congestion. Congestion is determined by receipt of BECN or ForeSight backward congestion notification messages. This adaptive shaping behavior was disabled in the configuration, so the net result of this configuration parameter is to satisfy CBWFQ’s calculation of its bandwidth (in kilobits per second) when classes are allocated using percentages.

For example, if using the following policy map, CBWFQ uses the MINCIR value, **frame-relay mincir 486400** to calculate the bandwidth in kilobits per second as illustrated in the subsequent **show policy map** example output:

```
policy-map llq-branch
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority percent 33
  class class-default
    fair-queue
```

From the following display, note that Cisco IOS software has calculated bandwidth for the call-setup class as 24 Kbps as 5 percent of the MINCIR value of 486400.

```
vpn18-2600-2#show policy-map interface s0/0.100
Serial0/0.100: DLCI 100 -

Service-policy output: llq-branch

Class-map: call-setup (match-all)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
Match: ip precedence 3
Weighted Fair Queueing
  Output Queue: Conversation 41
  Bandwidth 5 (%)
  Bandwidth 24 (kbps) Max Threshold 64 (packets)
  (pkts matched/bytes matched) 0/0
  (depth/total drops/no-buffer drops) 0/0/0
```

Attach Service Policy to Frame Relay Map Class

For branch routers after completing configuration of the QoS *Service-policy* verify it is attached to the Frame Relay **map-class**.

```
!
map-class frame-relay ts-branch
frame-relay cir 486400
frame-relay bc 4864
frame-relay be 0
frame-relay mincir 486400
no frame-relay adaptive-shaping
service-policy output llq-branch
frame-relay fragment 640
!
end
```

For a head-end WAN aggregation router using a VIP and Distributed Traffic Shaping (DTS), verify the shaper is attached to the Frame Relay **map-class**.

```
!
map-class frame-relay 512kb
no frame-relay adaptive-shaping
service-policy output 512kb-shaper
frame-relay fragment 640
!
end
```

Apply Traffic Shaping to the Output Interface

For branch routers enable **frame-relay traffic-shaping** to the physical interface and attach the Frame Relay **map-class** to all the subinterface DLCIs:

```
hostname vpn13-1700-4
!
interface Serial1/0
  encapsulation frame-relay
  frame-relay traffic-shaping
!
interface Serial1/0.1 point-to-point
  ip address 192.168.224.2 255.255.255.252
  frame-relay interface-dlci 101
  class ts-branch
end
```

The 7500 VIP configuration with Distributed Traffic Shaping at the head-end WAN aggregation router is configured similarly, and is shown below. Note that **frame-relay traffic-shaping** is not configured on the physical interface.

```
!
hostname vpn2-7500-2
!
interface Serial2/0/0/20:1
  description vpn13-1700-4
  no ip address
  encapsulation frame-relay
  no fair-queue
!
interface Serial2/0/0/20:1.102 point-to-point
  description vpn13-1700-4
  ip address 192.168.224.1 255.255.255.252
  frame-relay interface-dlci 102
  class 512kb
!
end
```

Applying Service Policy to HDLC Encapsulated T1 Interfaces

For implementations with T1 interfaces and HDLC encapsulation the following configuration would be used. The voice class is configured for 504 Kbps, which accommodates nine G.729 calls at 56 Kbps per call. The service policy is simply applied to the main interface. The clock rate of the interface provides congestion feedback- no shaping is required in this configuration. No Layer-2 fragmentation (LFI/FRF.12) is required at T1 line rates.

```
!
hostname vpn11-2600-4
!
policy-map 1536kb
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority 504
  class class-default
    fair-queue
!
interface Serial0/0
  bandwidth 1536
  ip address 192.168.154.2 255.255.255.252
  service-policy output 1536kb
  crypto map static-map
!
end
```

For the 7500 WAN aggregation router servicing this branch site, the following configuration was used during testing. The policy map is the same as the above branch router configuration.

```
!
hostname vpn2-7500-1
!
controller T3 2/0/0
  clock source line
  cablelength 50
  [...]
  t1 10 channel-group 1 timeslots 1-24
  [...]
!
interface Serial2/0/0/10:1
  description vpn11-2600-4
  ip address 192.168.154.1 255.255.255.252
  service-policy output 1536kb
!
```

Combined WAN and IPsec/IP GRE Router Configuration—Cisco 7200 HDLC/HSSI

The following configuration example applies to a Cisco 7200VXR router that is functioning as a WAN attached head-end router with a HSSI interface and HDLC encapsulation. This router is also configured as the IPsec/IP GRE head-end router with a tunnel interface and a crypto map entry for each remote peer. The *voice* class priority value would be calculated by multiplying the maximum total number of concurrent calls expected to the remote routers times the bandwidth per call. The **mission-critical** and **call-setup** classes are specified in percentages.



Note

In this example, since the service policy is matching on ToS byte, not other fields of the IP header (such as port number, protocol, source/destination IP address), it is not necessary to add the **qos pre-classify** command. (The example is on a Cisco 7200VXR with the Cisco IOS software 12.1(9)E image. Refer to the “[QoS Pre-Classify](#)” section on page 4-12 for more information regarding applicability of the **qos pre-classify** feature.) The original packet’s ToS byte is copied to the IPsec encapsulated header and is visible to the output service policy.

```
!
hostname vpn3-7200-1
!
boot system flash disk0:c7200-ik2s-mz.121-9.E.bin
!
ip cef
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
!
policy-map 17408kb
  class mission-critical
    bandwidth percent 22
  class voice
    priority 5544
  class call-setup
    bandwidth percent 5
  class class-default
    fair-queue
!
crypto map static-map local-address Hssi3/0
crypto map static-map 1 ipsec-isakmp
  set peer 192.168.1.2
  set transform-set vpn-test
  match address vpn-static1
crypto map static-map 2 ipsec-isakmp
  set peer 192.168.2.2
  set transform-set vpn-test
  match address vpn-static2
!
```



Note

One map entry for each peer.

```

!
interface Tunnel1
  description vpn6-2600-1
  ip address 10.62.1.193 255.255.255.252
  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
  crypto map static-map
!

```



Note One tunnel for each peer.

```

!
interface Hssi3/0
  description Hssi3/0
  bandwidth 17408
  ip address 192.168.251.1 255.255.255.0
  load-interval 30
  service-policy output 17408kb
  hssi internal-clock
  serial restart-delay 0
  crypto map static-map
!
ip access-list extended vpn-static1
  permit gre host 192.168.251.1 host 192.168.1.2
!

```



Note One access-list for each tunnel/crypto peer

IKE and IPSec Configuration

This section addressing configuration of Internet Key Exchange (IKE) and IPSec. These topics are addressed in a series of sections:

- [Configure ISAKMP Policy and Pre-shared Keys, page 6-20](#)
- [Configure IPSec Local Address, page 6-20](#)
- [Configure IPSec Local Address, page 6-20](#)
- [Configure IPSec Transform-Set, page 6-21](#)
- [Configure Crypto Map, page 6-21](#)
- [Apply Crypto Map to Interfaces, page 6-22](#)
- [Configuring QoS Pre-Classify, page 6-23](#)

The following sample configurations illustrate the parameters used for the Internet Security Association and Key Management Protocol (ISAKMP) and IPSec security policy.

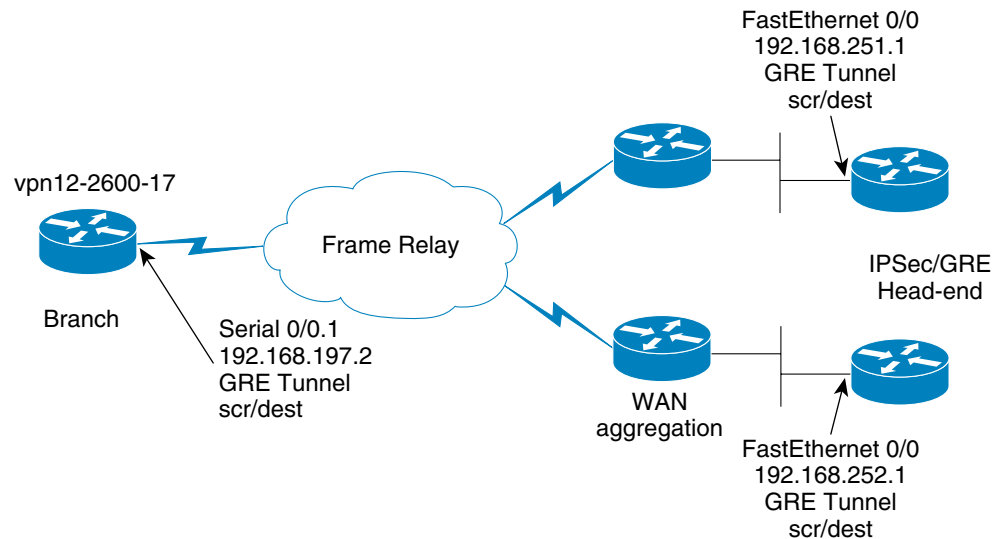
IKE is a hybrid protocol which implements the Oakley key exchange and Skeme key exchange inside the ISAKMP framework. ISAKMP, Oakley, and Skeme are security protocols implemented by IKE.

Security Associations (SA) are required by both IPSec and IKE. IKE negotiates and establishes its own SA and typically IPSec's SAs are created by IKE.

The IPSec security association lifetimes were the default values of 4608000 kilobytes/3600 seconds. The IKE lifetime was the default value of 24 hours, or 86,400 seconds.

Figure 6-3 illustrates the router addresses and interfaces for the following configuration examples.

Figure 6-3 IKE and IPSec Sample Topology



Configure ISAKMP Policy and Pre-shared Keys

The ISAKMP policy is configured to use group 2 (1024-bit Diffie-Hellman group.) Diffie-Hellman is a public-key protocol to establish session keys, a shared secret, over an unsecured path. Group 1, 768-bit Diffie-Hellman is also supported.

Pre-shared keys were used in lab testing. Pre-shared keys are commonly implemented, estimated at more than 75 percent of implementations—but present scalability challenges.

The IP addresses on the **crypto isakmp key** statements are the same addresses as configured subsequently as the **set peer** IP addresses in this router's crypto map. This statement also matches the head-end router's **crypto map map-name local-address interface-id** statement. There are two keys defined, one for each IPSec/IP GRE head-end router.

This is the branch router's ISAKMP policy and pre-shared key configuration.

```
!
hostname vpn12-2600-17
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.252.1
crypto isakmp key bigsecret address 192.168.251.1
!
end
```

This is a sample from one of the two head-end IPSec/IP GRE routers, in this case the router with the IP address of 192.168.252.1. Each head-end router has a key for each remote peer router's **crypto map map-name local-address interface-id** configuration command.

```
!
hostname vpn3-7200-2
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.1.2
crypto isakmp key bigsecret address 192.168.2.2
crypto isakmp key bigsecret address 192.168.3.2
...
crypto isakmp key bigsecret address 192.168.197.2
...
crypto isakmp key bigsecret address 192.168.244.2
!
end
```

Configure IPSec Local Address

Use of the **crypto map map-name local-address interface-id** can reduce overhead and makes administration easier. It allows the network administrator to determine the interface (and therefore the associated IP address) to be used as the router's identify to the remote peers. Use of the **local-address** can reduce the number of IKE security associations between two peers if they have multiple interfaces with crypto maps applied.

For routers with multiple or redundant interfaces (multiple paths to reach the IKE/IPsec peers) Loopback addresses are a best practice. This is a similar concept to defining Loopback interfaces and referencing the associated IP address in a **dls** **local-peer** configuration statement, or **snmp-server trap-source interface** statement.

In the case of the sample configuration, there is only one interface connecting to the WAN cloud, the example shows using the Serial interface for the branch.

```
!
hostname vpn12-2600-17
!
crypto map static-map local-address Serial0/0.1
!
```

The head-end example references the Fast Ethernet interface.

```
!
hostname vpn3-7200-2
!
crypto map static-map local-address FastEthernet0/0
!
```

In this example, only one **local-address** statement is needed per router.

Configure IPsec Transform-Set

This design guide implements Triple DES (168-bit /112-bit effective) rather than DES (56-bit). Cisco IOS software with strong encryption is subject to United States government export controls. Triple DES can have limited distribution and therefore might not be an option for use by all organization in all geographies. In general, the stronger the encryption the more computationally intensive. The lab testing represents the worst-case scenario.

SHA-1 is the hash algorithm (for authentication) used by both ISAKMP and IPsec. SHA-1 generates 20-byte hashes. The alternative, MD5—which generates 16-byte hashes—is not recommended as it is considered to have weaknesses. Both hash algorithms are truncated to 12 bytes in the ESP packet as described in RFC2104. The receiver computes the entire 20-byte value and compares the first 12 bytes with the value in the ESP packet

The following configuration is used in all branch and head-end routers.

```
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
```

It is common to see configurations that include both **ah-sha-hmac** and **esp-sha-hmac**. This is a duplication of the hashing or message integrity function and serves to increase the IPsec overhead. It is not recommended for this design.

Configure Crypto Map

The crypto map configuration ties together the IPsec components configured previously in this chapter.

In the following example, the **permit gre host** entries in the crypto access lists are the GRE tunnel interface's **tunnel source** and **tunnel destination** IP addresses. This **match address** statement defines what packets are being encrypted and authenticated by IPsec. In this design, that is the IP GRE tunnel that encapsulates the voice and data traffic.

The **set peer** statements reference the remote router's **local-address**, and the head-end routers reference this remote router's **local-address** in their crypto maps. There are two **crypto map** entries—sequence number 10 and 20—and a transmit and receive IPSec Security Association is created to each head-end router. This example is similar for the head-end routers, however, the map is repeated with the sequence number incremented. The head-end routers include a sequenced entry for each remote peer and an access list for each remote peer.

```
!
hostname vpn12-2600-17
!
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.252.1
  set transform-set vpn-test
  match address vpn-static1
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address vpn-static2
!
ip access-list extended vpn-static1
  permit gre host 192.168.197.2 host 192.168.252.1
ip access-list extended vpn-static2
  permit gre host 192.168.197.2 host 192.168.251.1
!
end
```

Since IP GRE tunnels are encrypted and EIGRP is configured for the tunnel interfaces, the EIGRP hello packets force IKE to continually build new IPSec SAs to transmit these hellos—even if no voice or user data traffic is being transmitted. By default, new IPSec SAs are created once per 3600 seconds (one per hour). Thus the IPSec tunnels are always up and available.

Apply Crypto Map to Interfaces

The **crypto map** configuration command *must* be applied to both the IP GRE tunnel interface and to the physical interface. For example:

```
!
hostname vpn12-2600-17
!
interface Tunnel0
  ip address 10.63.37.194 255.255.255.252
  tunnel source 192.168.197.2
  tunnel destination 192.168.252.1
  crypto map static-map
!
! Note, Tunnel 1 not shown
!
interface Serial0/0.1 point-to-point
  ip address 192.168.197.2 255.255.255.252
  frame-relay interface-dlci 101
  class ts-branch
  crypto map static-map
!
!
```

Configuring QoS Pre-Classify

When configuring an IPsec encrypted IP GRE tunnel enable **qos pre-classify** on both the Tunnel interface and crypto map. QoS Pre-Classify is not enabled by default in the releases tested.

```
!
class-map match-all call-setup
  match ip precedence 3
class-map match-all mission-critical
  match ip precedence 2
class-map match-all voice
  match ip precedence 5
!
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address vpn-static1
  qos pre-classify
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.252.1
  set transform-set vpn-test
  match address vpn-static2
  qos pre-classify
!
interface Tunnel0
  description Tunnel0
  ip address 10.62.139.194 255.255.255.252
  qos pre-classify
  tunnel source 192.168.91.2
  tunnel destination 192.168.251.1
  crypto map static-map
!
interface Tunnel1
  description Tunnel1
  ip address 10.62.139.198 255.255.255.252
  qos pre-classify
  delay 60000
  tunnel source 192.168.91.2
  tunnel destination 192.168.252.1
  crypto map static-map
!
end
```

**Note**

Cisco recommends QoS Pre-Classify be enabled on all branch VPN routers which support the feature.

Implementation and Configuration Checklist

This implementation and configuration checklist was compiled to help organize a successful implementation. Unless otherwise noted, these implementation steps must be similarly configured on both branch and head-end routers.

Table 6-3 Summary of Implementation Tasks

Implementation/Configuration Step	Section References
Configure IP GRE tunnel interfaces.	“Configure IP GRE Tunnels” section on page 6-2
Apply EIGRP manual summarization, ip summary-address to IP GRE tunnel interfaces.	“EIGRP Summarization and Network Addressing” section on page 6-2
Modify EIGRP hold-time if necessary.	“EIGRP hold-time” section on page 6-3
Increase delay value for backup IP GRE tunnel interface.	“IP GRE Tunnel Delay” section on page 6-3
Verify IP GRE tunnel interfaces are up/up and EIGRP neighbors are established.	“Verifying Tunnel Interfaces and EIGRP Neighbors” section on page 7-3
Verify campus switches/workstations/application /IP Phones are setting ToS byte accordingly.	“Using NetFlow to Verify ToS Values” section on page 7-6
Configure campus edge routers to map ToS to CoS—if applicable	“Campus QoS—Mapping ToS to CoS” section on page 6-5
Configure class-maps eg, voice, call-setup, mission-critical	“Configure QoS Class Map” section on page 6-6
Configure policy-map for WAN edge routers to allocate bandwidth for LLQ and percent classes.	“QoS Policy Map Configuration” section on page 6-7
Configure Frame Relay map-class with traffic shaping, and appropriate LFI per link speed.	“Frame Relay Traffic Shaping and FRF.12 (LFI)” section on page 6-11
Apply service-policy to Frame Relay map-class .	“Attach Service Policy to Frame Relay Map Class” section on page 6-14
Apply Frame Relay traffic-shaping to main interface, apply FRTS map-class to subinterface.	“Apply Traffic Shaping to the Output Interface” section on page 6-15
Apply Service Policy to T1 interfaces.	“Applying Service Policy to HDLC Encapsulated T1 Interfaces” section on page 6-16
Configure isakmp policy and pre-shared keys.	“Configure ISAKMP Policy and Pre-shared Keys” section on page 6-20
Configure ipsec local-address .	“Configure IPSec Local Address” section on page 6-20
Configure ipsec transform-set .	“Configure IPSec Transform-Set” section on page 6-21
Configure crypto map .	“Configure Crypto Map” section on page 6-21

Table 6-3 Summary of Implementation Tasks

Implementation/Configuration Step	Section References
Apply crypto map to Interfaces.	“Apply Crypto Map to Interfaces” section on page 6-22
Apply qos pre-classify .	“Configuring QoS Pre-Classify” section on page 6-23
Display IKE and IPSec configuration.	“Sample Show Commands for IPSec” section on page 7-8
Verify encrypting routers are not Layer-3 fragmenting packets.	“Packet Fragmentation” section on page 7-1



Verification and Troubleshooting

This chapter provides tips to assist in verification and troubleshooting the implementation. Specific troubleshooting discussions address the following:

- [Packet Fragmentation, page 7-1](#)
- [Displaying Anti-Replay Drops, page 7-2](#)
- [Verifying Tunnel Interfaces and EIGRP Neighbors, page 7-3](#)
- [How EIGRP calculates RTO values for Tunnel Interfaces, page 7-4](#)
- [Using NetFlow to Verify Layer-3 Packet Sizes, page 7-5](#)
- [Using NetFlow to Verify ToS Values, page 7-6](#)
- [Sample Show Commands for IPSec, page 7-8](#)
- [Clearing IPSec and IKE Security Associations, page 7-10](#)
- [Sample Show Commands for QoS, page 7-12](#)

Packet Fragmentation

IPSec and IP GRE headers increase the size of the original packet. [Chapter 4, “Planning and Design,”](#) illustrates how a 60-byte G.729 voice packet expands to 136 bytes after addition of the additional headers and trailer. While these relatively small voice packets would not exceed an interface’s MTU, data packets at or near MTU size could require fragmentation by the initial or intermediate routers.

Packet fragmentation should be avoided as it decreases router performance, both on the fragmenting router and by the end station. Since encryption is being done with IPSec routers, the end station could be the decrypting router. Fragmentation is performed after encryption and before decrypting; the decrypting router must process switch the packet since it must receive and re-assemble all fragments before decryption.

Fragmentation should be avoided by using either path MTU discovery or manually setting the MTU of the workstations to 1400 bytes. Cisco’s VPN Client installation provides a utility that changes the workstation’s MTU. From the Window’s task bar, select Start, Search, for Files or Folders and search for `setMTU.exe`. Execute this program, set the MTU to 1400 bytes and reboot.

To eliminate DLSw induced fragmentation consider defining a MAXDATA value which is smaller than the DLSw, IPsec and GRE overhead. The MAXDATA value is defined under the PU2.0 definition for the switched major node. This value indicates the maximum number of bytes a PU 2.0 device can send/receive. The value specified includes SNA overhead. For example:

```
MAXDATA=1033,
```

The default value (line 382) for a Cisco 3174 configuration is 521, a value of 265 is also commonly used.

To set the MTU on a Macintosh with OS X with the terminal program—it requires *sudo* or *root* access. Sudo access can be enabled through the *NetInfo Manager* application located under *Applications -> Utilities*. Users must go to the *Domain* pull down menu and under *Security* select *Enable Root Account*.

Step 1 Identify your network port with Terminal program: `ifconfig -a`

Step 2 Enter this command: `sudo /sbin/ifconfig en0 mtu 1400`



Note Assumes `en0` was the interface identified in prior step.

To display if the encrypting router is fragmenting packets, issue the following command several times while the network is in use:

```
sh ip traffic | include fragmented
4003204 fragmented, 0 couldn't fragment
```

If the fragmented counter is increasing, fragmentation is occurring. Refer to the [“Using NetFlow to Verify Layer-3 Packet Sizes”](#) section on page 7-5 for information about how to use NetFlow to verify packet sizes. Enable NetFlow switching on the interface shared with the workstations to determine the size of the packets prior to encryption.

Fragmentation does not degrade performance of intermediate routers not involved in the fragmentation or re-assembly process. Fragmented packets maintain the ToS byte of the original packet and intermediate router’s QoS policy is not affected.

Displaying Anti-Replay Drops

The procedure for displaying packets dropped due to the anti-replay logic differs depending if a hardware crypto accelerator is used or if encryption/decryption is done by software. Since hardware crypto accelerators are recommended for voice, those display examples are presented in this section. With hardware crypto accelerators, the sequence failures are checked and reported by the card, and are not IPsec Security Association (SA) specific, as is the case with software. The counters are an accumulation for all IPsec peers for this router:

For the Cisco 1700, Cisco 2600, Cisco 3600, and Cisco 3700 platforms the command is:

```
vpn13-1700-4#show crypto engine accelerator statistic | include esp_seq_fail
esp_prot_absent:0 esp_seq_fail: 0 esp_spi_failure: 0
```

For Cisco 7100 and Cisco 7200 platforms the commands are:

```
vpn3-7200-1#show pas isa int | include esp_seq_fail
esp_seq_failure: 249088 esp_spi_failure: 0
```

```
vpn3-7200-1#show pas vam int | include pkt_replay_err
rng_st_fail : 0 pkt_replay_err : 0
```


The counters are accumulated since the hardware crypto accelerator was last initialized or manually cleared with the **clear crypto engine accelerator counter** command.

Verifying Tunnel Interfaces and EIGRP Neighbors

Before attempting to configure IPsec and encrypt voice packets on the network, verify all the configured interfaces are in UP/UP state:

```
vpn13-1700-4#show interface | include is up
```

```
FastEthernet0/0 is up, line protocol is up
Serial1/0 is up, line protocol is up
Serial1/0.1 is up, line protocol is up
Loopback0 is up, line protocol is up
Tunnel0 is up, line protocol is up
Tunnel1 is up, line protocol is up
```

Verify that the IPsec/IP GRE head-end routers are neighbors over the tunnel interfaces. This display is from a branch router.

```
vpn13-1700-4#show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num	Type
1	10.63.88.193	Tu0	13 1w5d	47	5000	0	208349	
0	10.63.88.197	Tu1	13 1w5d	46	5000	0	302665	

For the same branch router, look at the routing table:

```
vpn13-1700-4#show ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route
```

```
Gateway of last resort is 192.168.224.1 to network 0.0.0.0
 192.168.224.0/30 is subnetted, 1 subnets
 C    192.168.224.0 is directly connected, Serial1/0.1
 10.0.0.0/8 is variably subnetted, 7 subnets, 5 masks
 D    10.0.0.0/8 [90/297246976] via 10.63.88.193, 1w5d, Tunnel0
 D    10.63.88.0/24 is a summary, 1w5d, Null0
 C    10.63.88.0/25 is directly connected, FastEthernet0/0
 C    10.63.88.254/32 is directly connected, Loopback0
 C    10.63.88.196/30 is directly connected, Tunnel1
 C    10.63.88.192/30 is directly connected, Tunnel0
 S*   0.0.0.0/0 [1/0] via 192.168.224.1
```

Note from the above display, network 10.0.0.0/8 was learned from the primary tunnel, Tunnel0. Only the route from the primary tunnel is inserted into the routing table, since in the configuration the interface delay for Tunnel1 was increased, making it an alternate or backup path. Also, note the summary route for 10.63.88.0/24 to the Null0 interface. This is the result of the manual summarization statement on the tunnel interfaces. In this example, only one EIGRP route is being learned through the Tunnel interfaces and only one route is being advertised to each IPsec/IP GRE head-end router.

How EIGRP calculates RTO values for Tunnel Interfaces

This design illustrates the use of EIGRP and GRE Tunnel interfaces. The default bandwidth for a Tunnel interface in Cisco IOS software is 9 Kbps. EIGRP calculates for each neighbor a Retransmission timeout (RTO) value in milliseconds. The RTO value is the amount of time Cisco IOS software waits before a retransmit of a reliable packet (EIGRP an update, query, reply) to its neighbor if an acknowledgement is not received.

The RTO value is computed by calculating:

- The current SRRT for the peer multiplied by 6
- The Pacing Timer for the interface multiplied by 6

Select the higher value of these two calculations. If either computed value is greater than 5,000 milliseconds, the RTO value is set to 5,000 milliseconds, or 5 seconds.

The pacing timer for an interface is calculated from the bandwidth value of the interface. The lower the bandwidth value, the higher the computed pacing timer. The pacing timer is the means to throttle EIGRP's utilization of an interface for routing protocol traffic. By default EIGRP uses up to 50 percent of the bandwidth available. This value can be changed by invoking the **ip bandwidth-percent eigrp** configuration command.

As an illustration, with the default tunnel bandwidth value of 9 Kbps, the RTO calculation would be $2702 * 6 = 16,212$ which is greater than 5,000, so the value of 5,000 is used for the RTO.

```
vpn-jk-2600-25#show interfaces tunnel 0 | include BW
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
```

```
vpn-jk-2600-25#show ip eigrp interfaces
IP-EIGRP interfaces for process 45
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/1	0	0/0	0	0/10	0	0
Tu0	1	0/0	649	71/2702	5894	0

```
vpn-jk-2600-25#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 45
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num	Tye
0	10.248.0.2	Tu0	13	05:05:26	649	5000	0	2	

Version 12.2/1.2, Retrans: 0, Retries: 0

Looking at a tunnel interface which was configured to use a bandwidth value of 56 Kbps, the RTO calculation would be: $434 * 6 = 2,604$

```
vpn-jk-2600-25#show interfaces tunnel 0 | include BW
  MTU 1514 bytes, BW 56 Kbit, DLY 500000 usec,
```

```
vpn-jk-2600-25#show ip eigrp interfaces
IP-EIGRP interfaces for process 45
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/1	0	0/0	0	0/10	0	0
Tu0	1	0/0	56	11/434	434	0

```
vpn-jk-2600-25#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 45
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num	Tye
0	10.248.0.2	Tu0	14	00:00:18	56	2604	0	4	

Version 12.2/1.2, Retrans: 1, Retries: 0


In either of the above two examples, the SRRT value multiplied by 6 is less than the RTO value multiplied by 6.

A RTO value of 5,000 in itself does not present a problem to the design. The use of manual summarization (and EIGRP stub) in this design minimizes the number of EIGRP updates, queries and replies which must be transmitted between branches and head-end routers. Increasing the bandwidth value for a tunnel interface decreases the pacing time, which allows EIGRP updates, queries and replies to be sent more frequently, but with good summarization the number of these transactions should be minimal.

Using NetFlow to Verify Layer-3 Packet Sizes

The topology shown in [Figure 7-1](#) is used in this section to illustrate the use of NetFlow to verify Layer-3 packet sizes.

Figure 7-1 Netflow Example Topology



Generate **60-byte** (Layer-3) packets with a traffic generator, then use the following command to capture packet information

```
vpn18-2 (TGN:ON, Fa0/1:1/1) #show ip
```

```
Summary of IP traffic streams on FastEthernet0/1
```

ts#	tos	len	id	frag	tll	protocol	chksm	source	destination
1	UDP	A0	60	0000	0000	60	17	6890 10.0.1.2	10.127.0.1

Given the following configuration of the router decrypting the traffic, verify the packet sizes by enabling Netflow on the Serial and Tunnel interfaces so the same flow is captured both encrypted and unencrypted:

```
!
hostname vpn18-2600-3
!
interface Tunnel0
 ip address 10.0.96.2 255.255.255.0
 ip route-cache flow
 tunnel source Loopback0
 tunnel destination 192.168.2.1
 crypto map GRE
!
interface Serial0/1
 no ip address
 encapsulation frame-relay
 ip route-cache flow
 frame-relay traffic-shaping
!
interface Serial0/1.100 point-to-point
 bandwidth 64
```

```

ip address 192.168.1.1 255.255.255.252
frame-relay interface-dlci 100
  class ts-branch
crypto map GRE
!

```

```

vpn18-2600-3#sh ip cache verbose flow | begin TOS
SrcIf          SrcIPAddress  DstIf          DstIPAddress  Pr TOS Flgs  Pkts
Port Msk AS      Port Msk AS    NextHop        B/Pk Active
Tu0            10.0.1.2      Fa0/1          10.127.0.1    11 A0 10    3607
7D05 /0 0      7D09 /24 0    10.254.0.45   60    72.1
Se0/1.100     192.168.1.1  Local          192.168.1.2  32 00 10    2462
E74D /30 0     2454 /30 0    0.0.0.0       136   49.1
Tu0           192.168.2.1  Null           192.168.1.6  11 00 10
17 01F4 /0 0   01F4 /0 0    0.0.0.0       112  140.9

```

In the **show ip cache** output, the first flow is the UDP packets from the traffic generator, after they were decrypted; the second line shows the IPsec ESP packet (protocol 50) from the Serial interface before it was decrypted; and the last packet is an IKE packet. The increase in packet size (NetFlow reports Layer-3 packet lengths) can be calculated by subtracting the average bytes per packet before and after the IPsec and IP GRE headers. In this case the original packet was 60 bytes, with IPsec (tunnel mode) and IP GRE 136 bytes.

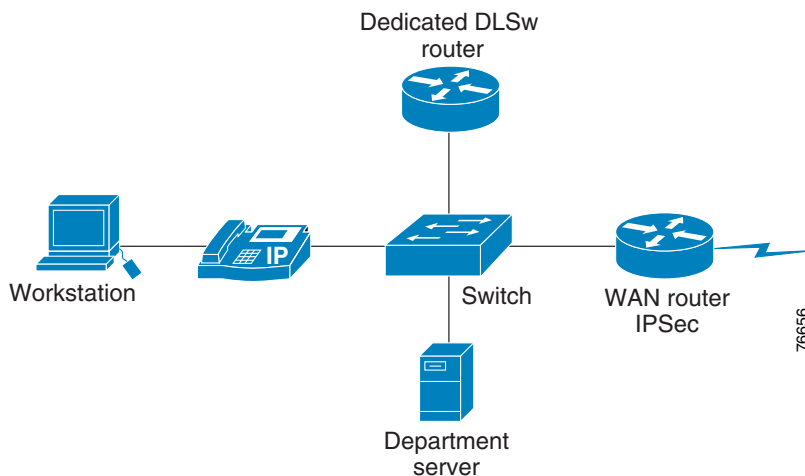
**Note**

In the preceding **show ip cache** output, NetFlow reports the ToS byte as zero. NetFlow maps the bits from the *ip_more_fragment* flag into the ToS byte for IPsec ESP (protocol 50) and AH (protocol 51) tunnels. IP GRE tunnels are handled differently. IP GRE flows are displayed as separate flows if the pre-IP GRE tunnel-encapsulated packets have varying ToS bytes.

Using NetFlow to Verify ToS Values

The topology shown below in [Figure 7-2](#) is used to illustrate using NetFlow to verify ToS values from the LAN. There is a dedicated DLSw router advertising its loopback interface via EIGRP to the WAN router that would be the IPsec peer router.

Figure 7-2 Netflow ToS Verification Topology



This is a ToS verification technique which can be done remotely or without the need for a protocol analyzer on the LAN/WAN. This is an alternative to exporting NetFlow to a Collector/Analyzer or other third party collection device. On the WAN router enable NetFlow switching:

```
!
interface FastEthernet0/1
 ip address 10.254.0.45 255.255.255.0
 ip route-cache flow
end

vpn-18-2600-5#show ip cache verbose flow | begin TOS
SrcIf          SrcIPAddress      DstIf          DstIPAddress    Pr TOS Flgs  Pkts
Port Msk AS    Port Msk AS      NextHop        B/Pk Active
Fa0/1          10.254.0.47        Null           224.0.0.10      58 C0 10    116
0000 /0 0      0000 /0 0        0.0.0.0         60 533.1
Fa0/1          10.251.0.1        Local          10.254.0.45    06 40 18     2
0811 /0 0      2B06 /0 0        0.0.0.0         46 0.0
```

From the above **show ip cache** output, note an EIGRP hello and DLSw flow. These values need to be converted from hex to decimal.

The EIGRP flow is identified by protocol 88 (0x58).



Note

The TOS byte is 0xC0, or IP Precedence 6. See [Figure 7-3](#).

DLSw listens by default on TCP (protocol 0x06) port 2065 (0x0811). In this example, the default IP Precedence value for DLSw was changed from 5 to 2, to match the service policy's "mission-critical" class.

Figure 7-3 ToS Byte Anatomy




Table 7-1 ToS Byte Reference

TOS Hex	TOS Decimal ¹	IP Precedence	Class-map Name (Used in Examples)	DSCP	Binary
E0	224	7 Network Control		CS7	111 00000
C0	192	6 Internetwork Control	mission-critical	CS6	110 00000
B8	184		voice	EF	101 11000
A0	160	5 Critical		CS5	101 00000
80	128	4 Flash Override		CS4	100 00000

Table 7-1 ToS Byte Reference

TOS Hex	TOS Decimal ¹	IP Precedence	Class-map Name (Used in Examples)	DSCP	Binary
68	104		call-setup	AF31	011 01000
60	96	3 Flash		CS3	011 00000
40	64	2 Immediate	mission-critical	CS2	010 00000
20	32	1 Priority		CS1	001 00000
00	0	0 Routine		default	000 00000

1. If the TOS hex value converted to decimal falls between the illustrated decimal values, IP Precedence matches on the next lower value. For example, TOS decimal values between 160 and 191 are IP Precedence 5.

The relevant configuration is as follows:

```
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
!
```

Sample Show Commands for IPsec

The following commands can be used to verify the implementation values are consistent with these recommendations:

The **show crypto map** command is used to verify the crypto map configuration. This command facilitates verification of the crypto local and peer IP addresses and the GRE IP addresses match, since in the configuration this information is normally not visible on one page. Also helps to verify that the crypto map is applied to the appropriate output and tunnel interfaces.

```
vpn13-1700-4#show crypto map
Crypto Map: "static-map" idb: Serial1/0.1 local address: 192.168.224.2

Crypto Map "static-map" 10 ipsec-isakmp
Peer = 192.168.252.1
Extended IP access list vpn-static1
  access-list vpn-static1 permit gre host 192.168.224.2 host 192.168.252.1
Current peer: 192.168.252.1
Security association lifetime: 4608000 kilobytes/3600 seconds
PFS (Y/N): N
Transform sets={
vpn-test,
}

Crypto Map "static-map" 20 ipsec-isakmp
Peer = 192.168.251.1
Extended IP access list vpn-static2
  access-list vpn-static2 permit gre host 192.168.224.2 host 192.168.251.1
Current peer: 192.168.251.1
```

```

Security association lifetime: 4608000 kilobytes/3600 seconds
PFS (Y/N): N
Transform sets={
vpn-test,
}
Interfaces using crypto map static-map:
Serial1/0.1
Tunnel0
Tunnel1

```

With the **show crypto isakmp sa** command, the state is normally `QM_IDLE` and there are two IKE security associations, one to each head end router, from the branch perspective. Either the branch or the head-end router can initiate an IKE session, so the destination (dst) and source (src) IP addresses don't have any particular meaning or affinity.

```

vpn13-1700-4#show crypto isakmp sa
dst          src          state         conn-id    slot
192.168.224.2 192.168.252.1 QM_IDLE      1          0
192.168.251.1 192.168.224.2 QM_IDLE      2          0

```

The **show crypto engine connections active** shows both IKE Security associations as well as IPsec. From the previous display, the connection-id of the IKE SAs are 1 and 2, and they are both shown below. From the branch router's perspective, there should normally be four IPsec SAs, since there are two head-end routers and there is a transmit (Encrypt) and receive (Decrypt) SA for each head-end. Note that since the recommended configuration has a primary and secondary IP GRE tunnel, the packet counts are much higher to the primary head-end than to the secondary head-end. Only EIGRP hellos and any other background traffic are traversing the tunnel to the secondary head-end unless the primary fails.

```

vpn13-1700-4#show crypto engine connections active

```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	Se1/0.1	192.168.224.2	set	HMAC_SHA+3DES_56_C	0	0
2	Tunnel0	10.63.88.194	set	HMAC_SHA+3DES_56_C	0	0
1910	Tunnel0	10.63.88.194	set	HMAC_SHA+3DES_56_C	0	567
1911	Tunnel0	10.63.88.194	set	HMAC_SHA+3DES_56_C	567	0
1912	Tunnel0	10.63.88.194	set	HMAC_SHA+3DES_56_C	0	72
1913	Tunnel0	10.63.88.194	set	HMAC_SHA+3DES_56_C	72	0

Use the **show crypto isakmp policy** to verify the IKE policy and how it differs from the default configuration.

```

vpn13-1700-4#show crypto isakmp policy
Protection suite of priority 1
encryption algorithm: Three key triple DES
hash algorithm: Secure Hash Standard
authentication method: Pre-Shared Key
Diffie-Hellman group: #2 (1024 bit)
lifetime: 86400 seconds, no volume limit
Default protection suite
encryption algorithm: DES - Data Encryption Standard (56 bit keys).
hash algorithm: Secure Hash Standard
authentication method: Rivest-Shamir-Adleman Signature
Diffie-Hellman group: #1 (768 bit)
lifetime: 86400 seconds, no volume limit

```

Use the **show crypto ipsec transform-set** to verify of the transform set options as well as tunnel verses transport mode for IPsec.

```

vpn13-1700-4#show crypto ipsec transform-set
Transform set vpn-test: { esp-3des esp-sha-hmac }
negotiate = { Tunnel, },

```

Clearing IPsec and IKE Security Associations

When making configuration changes or after a router reload, security associations (SAs) can become 'stale', (invalid or out of sync) between two IPsec peers. In some instances this can be the cause for IPsec connectivity failures. This is more common when using IPsec without GRE tunnels and a routing protocol, as the data traffic of the hello packets from the routing protocol forces new SAs to be built to receive and transmit the hellos.

Two clear commands can be used to flush the database and eliminate any legacy negotiations:

```
clear crypto isakmp
```

```
clear crypto sa
```

Here is an example to illustrate this point and steps through clearing both the IKE and IPsec SAs. Router *vpn18-2600-22* is a head-end IPsec/GRE router with one remote router, *vpn18-2600-18*, which has an IPsec/GRE tunnel to two head-end routers.

From the *vpn18-2600-22* device's perspective, there is an IKE and a transmit and receive IPsec SA to the remote router. The remote router is reloaded to simulate a branch failure; note that the EIGRP neighbor goes down and returns.

```
vpn18-2600-22#show crypto engine connections active
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	FastEthernet0/1	23.0.1.22	set	HMAC_SHA+3DES_56_C	0	0
2426	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	0	312
2427	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	312	0

```
vpn18-2600-22#
```

```
1w4d: %DUAL-5-NBRCHANGE: IP-EIGRP 45: Neighbor 10.96.1.1 (Tunnel0) is down: hold
```

```
vpn18-2600-22#
```

```
1w4d: %DUAL-5-NBRCHANGE: IP-EIGRP 45: Neighbor 10.96.1.1 (Tunnel0) is up: new ay
```

```
vpn18-2600-22#show crypto engine connections active
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	FastEthernet0/1	23.0.1.22	set	HMAC_SHA+3DES_56_C	0	0
2	FastEthernet0/1	23.0.1.22	set	HMAC_SHA+3DES_56_C	0	0
2428	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	0	0
2429	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	0	0
2430	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	0	154
2431	Tunnel0	10.96.1.2	set	HMAC_SHA+3DES_56_C	154	0



Note

Connection ID 2428 and 2429 are extraneous, they are not being used to encrypt or decrypt traffic. From the branch perspective, following the reload.

```
vpn18-2600-18#show crypto engine connections active
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
2	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
420	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	0
421	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	0
422	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	187
423	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	185	0
424	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	0
425	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	0
426	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	184
427	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	184	0

The branch has two active IPsec SAs to each head-end, Ids 422/423 and 426/427 and two IKE SAs, 1 and 2. Clearing the SAs eliminates the redundant SAs.

```
vpn18-2600-18#clear crypto sa
vpn18-2600-18#show crypto engine connections active
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
2	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
420	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	2
421	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	2	0
422	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	2
423	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	1	0

After the clear command, there are two IKE SAs, and four IPsec SAs, a transmit and receive tunnel to each head-end.

Now, looking at the IKE SAs, they are in a normal state Quick-Mode Idle, clearing them deletes the IKE SAs.

```
vpn18-2600-18#show crypto isakmp sa
dst          src          state          conn-id  slot
23.0.32.22  23.0.218.1  QM_IDLE       2        0
23.0.32.23  23.0.218.1  QM_IDLE       1        0
```

```
vpn18-2600-18#clear crypto isakmp
vpn18-2600-18#show crypto isakmp sa
dst          src          state          conn-id  slot
23.0.32.22  23.0.218.1  MM_NO_STATE   2        0  (deleted)
23.0.32.23  23.0.218.1  MM_NO_STATE   1        0  (deleted)
```

```
vpn18-2600-18#show crypto isakmp sa
dst          src          state          conn-id  slot
```



Note

The IKE SAs were deleted but have not been re-established, they are not needed at this time, since the IPsec SAs have not expired. They are still encrypting and decrypting packets. New IKE SAs are not required until the IPsec SAs timeout (exceed their lifetime, either triggered by time or data volume) and must be re-established.

```
vpn18-2600-18#show crypto engine connections act
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
420	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	63
421	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	62	0
422	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	62
423	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	61	0

To force IKE SAs establishment, clear the IPsec SAs (IKE SAs need to be build to establish new IPsec SAs).

```
vpn18-2600-18#clear crypto sa
vpn18-2600-18#show crypto engine connections act
```

ID	Interface	IP-Address	State	Algorithm	Encrypt	Decrypt
1	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
2	<none>	<none>	set	HMAC_SHA+3DES_56_C	0	0
420	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	11
421	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	12	0
422	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	0	11
423	Tunnel0	10.96.1.1	set	HMAC_SHA+3DES_56_C	11	0

```
vpn18-2600-18#show crypto isakmp sa
```

dst	src	state	conn-id	slot
23.0.32.22	23.0.218.1	QM_IDLE	1	0
23.0.32.23	23.0.218.1	QM_IDLE	2	0

From the above display, the router is functioning normally, and the expected number of security associations are seen in the display.

Sample Show Commands for QoS

Use the **show policy-map interface** to verify the offered rate of traffic isn't greater than the allocated bandwidth for the voice, call-setup and mission-critical classes, as drops in these classes impact voice quality, call setup and the important data applications.

If there are drops in the voice class, either the call admission control configuration is not consistent with the amount of bandwidth allocated for voice calls, or there could be a minor amount of jitter in the call that is causing the voice packets to arrive slightly over the rate calculated per call. If the call admission control issue was verified, the CODEC type isn't G.711 when it was planned to be G.729, and there is a minor amount of voice being dropped, then increase the value of the priority queue until the drops stop.

If there are drops in the IP Precedence 6 class within mission-critical, expect to experience EIGRP neighbors drop. The service policy should be tuned to prevent EIGRP packets from being dropped, as EIGRP packet drops cause instability in the network.

In the release of code tested, the *offered rate* (in bits per second) does not include GRE and IPsec header overhead; however, the packets per second rates should report accurate packet rates.

```
vpn13-1700-4#show policy-map interface serial 1/0.1
Serial1/0.1: DLCI 101 -

Service-policy output: llq-branch

Class-map: call-setup (match-all)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: ip precedence 3
  Weighted Fair Queueing
    Output Queue: Conversation 41
    Bandwidth 5 (%)
    Bandwidth 24 (kbps) Max Threshold 64 (packets)
    (pkts matched/bytes matched) 0/0
    (depth/total drops/no-buffer drops) 0/0/0

Class-map: mission-critical (match-any)
  2992 packets, 392844 bytes
```

```

30 second offered rate 2000 bps, drop rate 0 bps
Match: ip precedence 2
    0 packets, 0 bytes
    30 second rate 0 bps
Match: ip precedence 6
    2992 packets, 392844 bytes
    30 second rate 2000 bps
Weighted Fair Queueing
Output Queue: Conversation 42
Bandwidth 22 (%)
Bandwidth 106 (kbps) Max Threshold 64 (packets)
(pkts matched/bytes matched) 2994/393695
(depth/total drops/no-buffer drops) 0/0/0

Class-map: voice (match-all)
    0 packets, 0 bytes
    30 second offered rate 0 bps, drop rate 0 bps
Match: ip precedence 5
Weighted Fair Queueing
    Strict Priority
    Output Queue: Conversation 40
    Bandwidth 168 (kbps) Burst 4200 (Bytes)
    (pkts matched/bytes matched) 0/0
    (total drops/bytes drops) 0/0

Class-map: class-default (match-any)
    26601 packets, 10030161 bytes
    30 second offered rate 0 bps, drop rate 0 bps
Match: any
Weighted Fair Queueing
    Flow Based Fair Queueing
    Maximum Number of Hashed Queues 32
    (total queued/total drops/no-buffer drops) 0/0/0

```

Use the **show frame-relay fragment** command to verify the fragment size in bytes, and the number of packets that require fragmentation.

```

vpn13-1700-4#show frame-relay fragment
interface      dlci  frag-type  frag-size  in-frag  out-frag  dropped-frag
Serial1/0.1    101   end-to-end  640        154      154       0

```





Network Diagram Scalability Testbed and Configuration Files

This appendix contains configurations that were used during a V³PN performance and scalability evaluation based on the network illustrated in [Figure A-1](#). Specific configurations address the following devices and supporting networking functions:

- [Head-end VPN Router, page A-2](#)
- [Branch VPN Router—Frame Relay, page A-5](#)
- [Branch VPN Router—HDLC, page A-8](#)

Figure A-1 V³PN Solution Testbed Diagram



Head-end VPN Router

The configuration below was taken from the Cisco 7200 VPN Router being used as a head-end. In this configuration, QoS was enabled on a separate WAN aggregation device, not on the same router terminating VPN tunnels.

As the configuration is extremely large in its entirety (due to the repetition involved to configure all 244 branches being terminated), repetitive commands were removed and noted.

```

!
version 12.1
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
no service password-encryption
service compress-config
!
hostname vpn3-7200-2
!
boot system flash disk0:c7200-ik2s-mz.121-9.E.bin
logging buffered 65535 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
clock calendar-valid
ip subnet-zero
ip cef
!
no ip domain-lookup
!
ip ssh time-out 120
ip ssh authentication-retries 3
!
xsm
xsm privilege configuration level 15
xsm privilege monitor level 1
xsm vdm
xsm edm
no xsm history vdm
no xsm history edm
!
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.244.2
crypto isakmp key bigsecret address 192.168.242.2
.
<repetition removed>
.
crypto isakmp key bigsecret address 192.168.1.2
!
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
crypto mib ipsec flowmib history tunnel size 200
crypto mib ipsec flowmib history failure size 200
!
crypto map static-map local-address FastEthernet1/0
crypto map static-map 1 ipsec-isakmp
  set peer 192.168.1.2
  set security-association lifetime seconds 86400

```

```
    set transform-set vpn-test
    set pfs group2
    match address vpn-static1
crypto map static-map 2 ipsec-isakmp
    set peer 192.168.2.2
    set security-association lifetime seconds 86400
set transform-set vpn-test
    set pfs group2
    match address vpn-static2
.
<repetition removed>
.
crypto map static-map 244 ipsec-isakmp
    set peer 192.168.244.2
    set security-association lifetime seconds 86400
    set transform-set vpn-test
    set pfs group2
    match address vpn-static244
!
controller ISA 2/1
!
buffers small permanent 2048
buffers small max-free 10240
buffers small min-free 512
buffers middle permanent 2048
buffers middle max-free 10240
buffers middle min-free 512
buffers big permanent 2048
buffers big max-free 10240
buffers big min-free 512
buffers verybig permanent 2048
buffers verybig max-free 10240
buffers verybig min-free 512
buffers large permanent 2048
buffers large max-free 10240
buffers large min-free 512
buffers huge permanent 128
buffers huge max-free 512
buffers huge min-free 32
!
!
interface Loopback0
    description Loopback0
    ip address 10.57.2.255 255.255.255.255
!
interface Tunnel1
    description vpn6-2600-1
    ip address 10.62.1.197 255.255.255.252
    ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
    load-interval 30
    delay 60000
    tunnel source 192.168.252.1
    tunnel destination 192.168.1.2
    crypto map static-map
!
interface Tunnel2
    description vpn6-2600-2
    ip address 10.62.2.197 255.255.255.252
    ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
    load-interval 30
    delay 60000
    tunnel source 192.168.252.1
    tunnel destination 192.168.2.2
    crypto map static-map
```

```

.
<repetition removed>
.
interface Tunnel244
  description vpn17-4200-2
  ip address 10.63.130.193 255.255.255.252
  ip summary-address eigrp 1 10.0.0.0 255.0.0.0 5
  load-interval 30
  tunnel source 192.168.252.1
  tunnel destination 192.168.244.2
  crypto map static-map
!
interface FastEthernet0/0
  description FastEthernet0/0
  ip address 172.26.156.18 255.255.254.0
  load-interval 30
  duplex full
!
interface FastEthernet1/0
  description FastEthernet1/0
  ip address 192.168.252.1 255.255.255.0
  load-interval 30
  duplex full
  speed 100
  crypto map static-map
!
interface FastEthernet1/1
  description FastEthernet1/1
  ip address 10.57.2.1 255.255.255.252
  load-interval 30
  duplex full
  speed 100
!
interface Hssi3/0
  ip address 192.168.253.10 255.255.255.252
  shutdown
  hssi internal-clock
  serial restart-delay 0
!
router eigrp 1
  passive-interface FastEthernet0/0
  passive-interface FastEthernet1/0
  network 10.0.0.0
  no auto-summary
  eigrp log-neighbor-changes
!
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.252.2
ip route 172.26.0.0 255.255.0.0 172.26.156.1
no ip http server
!
!
ip access-list extended vpn-static1
  permit gre host 192.168.252.1 host 192.168.1.2
ip access-list extended vpn-static10
  permit gre host 192.168.252.1 host 192.168.10.2
ip access-list extended vpn-static100
  permit gre host 192.168.252.1 host 192.168.100.2
.
<repetition removed>
.
ip access-list extended vpn-static244
  permit gre host 192.168.252.1 host 192.168.244.2
logging trap debugging

```



```

logging 172.26.131.82
snmp-server community private RW
snmp-server community public RO
!
line con 0
  exec-timeout 0 0
  password cisco
  login
line aux 0
line vty 0 4
  exec-timeout 0 0
  password cisco
  login
line vty 5 15
  login
!
ntp clock-period 17179932
ntp server 172.26.156.1
end
!
```

Branch VPN Router—Frame Relay

The configuration shown below is from a Cisco 2651 VPN Router that was configured for V³PN. The Layer-2 technology used in this case was Frame Relay at a 1280 Kbps link speed.

```

!
! No configuration change since last restart
!
version 12.2
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
no service password-encryption
!
hostname vpn12-2600-1
!
logging buffered 32768 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
ip subnet-zero
ip cef
!
no ip domain-lookup
!
ip audit notify log
ip audit po max-events 100
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice
  match ip precedence 5
!
policy-map 1280kb
  class call-setup
    bandwidth percent 5
  class mission-critical
```

```

        bandwidth percent 22
    class voice
        priority 392
    class class-default
        fair-queue
!
crypto isakmp policy 1
    encr 3des
    authentication pre-share
    group 2
crypto isakmp key bigsecret address 192.168.252.1
crypto isakmp key bigsecret address 192.168.251.1
!
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto map static-map local-address Serial0/0.1
crypto map static-map 10 ipsec-isakmp
    set peer 192.168.252.1
    set transform-set vpn-test
    match address vpn-static1
crypto map static-map 20 ipsec-isakmp
    set peer 192.168.251.1
    set transform-set vpn-test
    match address vpn-static2
!
!
fax interface-type fax-mail
mta receive maximum-recipients 0
!
!
interface Loopback0
    ip address 10.63.21.254 255.255.255.255
!
interface Tunnel0
    description Tunnel0
    ip address 10.63.21.194 255.255.255.252
    ip summary-address eigrp 1 10.63.21.0 255.255.255.0 5
    load-interval 30
    qos pre-classify
    tunnel source 192.168.181.2
    tunnel destination 192.168.252.1
    crypto map static-map
!
interface Tunnel1
    description Tunnel1
    ip address 10.63.21.198 255.255.255.252
    ip summary-address eigrp 1 10.63.21.0 255.255.255.0 5
    load-interval 30
    delay 60000
    qos pre-classify
    tunnel source 192.168.181.2
    tunnel destination 192.168.251.1
    crypto map static-map
!
!
interface FastEthernet0/0
    description FastEthernet0/0
    ip address 172.26.157.181 255.255.254.0
    no ip proxy-arp
    no ip mroute-cache
    load-interval 30
    speed auto
    half-duplex

```

```
!
interface Serial0/0
  description Serial0/0
  bandwidth 1280
  no ip address
  encapsulation frame-relay
  no ip mroute-cache
  logging event subif-link-status
  logging event dlci-status-change
  load-interval 30
  no fair-queue
  frame-relay traffic-shaping
!
interface Serial0/0.1 point-to-point
  description Serial0/0.1
  bandwidth 1280
  ip address 192.168.181.2 255.255.255.252
  no ip mroute-cache
  frame-relay interface-dlci 101
  class 1280kb
  crypto map static-map
!
interface FastEthernet0/1
  description FastEthernet0/1
  ip address 10.63.21.1 255.255.255.128
  no ip mroute-cache
  load-interval 30
  speed 10
  full-duplex
!
router eigrp 1
  passive-interface Serial0/0
  passive-interface Serial0/0.1
  passive-interface FastEthernet0/1
  network 10.0.0.0
  no auto-summary
  eigrp log-neighbor-changes
!
ip default-gateway 192.168.181.1
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.181.1
ip route 10.63.21.200 255.255.255.255 10.63.21.2
no ip http server
ip pim bidir-enable
!
ip access-list extended vpn-static1
  permit gre host 192.168.181.2 host 192.168.252.1
ip access-list extended vpn-static2
  permit gre host 192.168.181.2 host 192.168.251.1
!
map-class frame-relay 1280kb
  no frame-relay adaptive-shaping
  frame-relay cir 1216000
  frame-relay bc 12160
  frame-relay be 0
  frame-relay mincir 1216000
  service-policy output 1280kb
!
snmp-server engineID local 000000090200000628DBD3E0
snmp-server community private RW
snmp-server community public RO
call rsvp-sync
!
mgcp profile default
```

```

!
dial-peer cor custom
!
line con 0
  exec-timeout 0 0
line aux 0
line vty 0 4
  exec-timeout 0 0
  password cisco
  login
  length 30
line vty 5 15
  login
!
ntp clock-period 17208540
ntp server 172.26.156.1
!
end
!

```

Branch VPN Router—HDLC

The configuration shown below is from a Cisco 1751 VPN Router that was configured for V³PN. The Layer-2 technology used in this case was HDLC at an E1 link speed.

```

!
! No configuration change since last restart
!
version 12.2
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
no service password-encryption
!
hostname vpn17-1700-1
!
logging buffered 65535 debugging
enable password cisco
!
clock timezone EST -5
clock summer-time EDT recurring
mmi polling-interval 60
no mmi auto-configure
no mmi pvc
mmi snmp-timeout 180
ip subnet-zero
!
!
no ip domain-lookup
!
ip audit notify log
ip audit po max-events 100
ip cef
ip ssh time-out 120
ip ssh authentication-retries 3
!
class-map match-all call-setup
  match ip precedence 3
class-map match-any mission-critical
  match ip precedence 2
  match ip precedence 6
class-map match-all voice

```

```
    match ip precedence 5
  !
  !
policy-map 2048kb
  class mission-critical
    bandwidth percent 22
  class voice
    priority 672
  class call-setup
    bandwidth percent 5
  class class-default
    fair-queue
  !
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp key bigsecret address 192.168.252.1
  !
  !
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
  !
crypto map static-map local-address Serial1/0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address vpn-static1
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.252.1
  set transform-set vpn-test
  match address vpn-static2
  !
  !
interface Loopback0
  ip address 10.63.100.254 255.255.255.255
  !
interface Tunnel0
  description Tunnel0
  ip address 10.63.100.198 255.255.255.252
  ip summary-address eigrp 1 10.63.100.0 255.255.255.0 5
  load-interval 30
  delay 60000
  qos pre-classify
  tunnel source 192.168.236.2
  tunnel destination 192.168.251.1
  crypto map static-map
  !
  !
interface Tunnel1
  description Tunnel1
  ip address 10.63.100.194 255.255.255.252
  ip summary-address eigrp 1 10.63.100.0 255.255.255.0 5
  load-interval 30
  qos pre-classify
  tunnel source 192.168.236.2
  tunnel destination 192.168.252.1
  crypto map static-map
  !
interface Ethernet0/0
  description FlashNet
  ip address 172.26.157.253 255.255.254.0
  half-duplex
  !
```

```

interface FastEthernet0/0
description FastEthernet0/0
ip address 10.63.100.1 255.255.255.128
load-interval 30
speed 100
full-duplex
!
interface Serial1/0
description Serial1/0
bandwidth 2048
ip address 192.168.236.2 255.255.255.252
no ip mroute-cache
load-interval 30
service-policy output 2048kb
crypto map static-map
!
router eigrp 1
network 10.0.0.0
no auto-summary
eigrp log-neighbor-changes
!
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.236.1
ip route 172.18.0.0 255.255.0.0 172.26.156.1
ip route 172.26.0.0 255.255.0.0 172.26.156.1
no ip http server
ip pim bidir-enable
!
!
ip access-list extended vpn-static1
permit gre host 192.168.236.2 host 192.168.251.1
ip access-list extended vpn-static2
permit gre host 192.168.236.2 host 192.168.252.1
!
!
snmp-server engineID local 0000000902000003E38D8C20
snmp-server community private RW
snmp-server community public RO
!
line con 0
exec-timeout 0 0
line aux 0
line vty 0 4
exec-timeout 0 0
password cisco
login
!
no scheduler allocate
ntp clock-period 17180765
ntp server 172.26.156.1
end
!

```



Configuration Supplement—Voice Module, EIGRP Stub, DSCP, HDLC

This appendix contains supplemental configurations used during a V³PN performance and scalability evaluation. Specific configurations address the following devices and supporting networking functions:

- [Voice Module Configuration, page B-1](#)
- [Router Configuration—vpn18-2600-2, page B-3](#)
- [Router Configuration—vpn18-2600-3, page B-4](#)
- [Router Configuration—vpn18-2600-4, page B-5](#)
- [Router Configuration—vpn18-2600-8, page B-6](#)
- [Router Configuration—vpn18-2600-9, page B-7](#)
- [Router Configuration—vpn18-2600-10, page B-8](#)
- [Router Configuration—vpn18-2600-6, page B-10](#)

Voice Module Configuration

The full-scale solution test was designed to validate a site-to-site VoIP over IPsec solution where the voice bearer traffic would be received on the LAN interface rather than generated locally by the router from a voice network module.

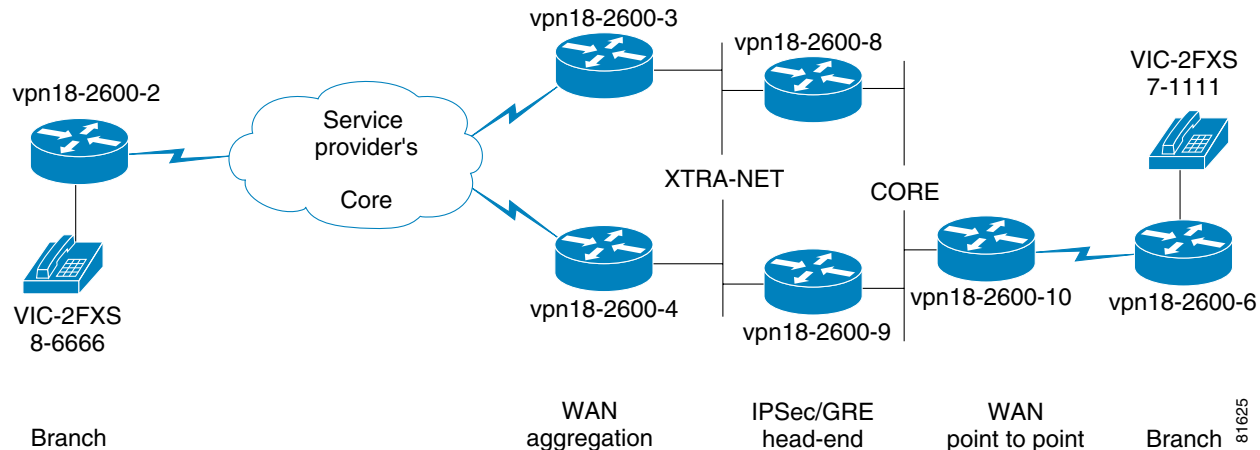
The purpose of this appendix is to create a single rack (small scale) configuration to supplement the main design guide to include the follow capabilities:

- IPsec *transport* mode configuration example
- EIGRP stub configuration
- Redundant configuration for dual WAN aggregation routers in addition to dual IPsec/GRE head-end routers
- Implement DSCP based *class-map*

HDLC sample configuration for a point-to-point WAN link

[Figure B-1](#) illustrates an example network configuration.

Figure B-1 Topology Diagram



These configuration examples do not include *class-map*, *policy-map* or *ISAKMP policy* configurations unless they differ from the configurations illustrated previously in this design guide.

To simplify the topology drawing, the interface descriptions in the following router configurations include the keywords *XTRA-NET* and *CORE* or in the case of serial links, the router on the opposite end of the link is listed. The addressing scheme is configured such that the 192.168.x.0 subnets would represent routable addresses (Non-RFC 1918) and the 10.0.0.0 address space would be representative of where an enterprise might deploy that address space.

The third octet of the loopback 0 interface on the devices shown is the same as the last digit of the host name. For example, router *vpn18-2600-8* has the loopback address of 192.168.8.1.

It should be noted, that while IPsec transport mode decreases the WAN interface bandwidth requirements, it does not decrease the number of packets per second, which in most cases, is the limiting factor of a router's performance. The **priority** keyword of the voice class in the policy-map was not decreased from the value used in the design guide—bandwidth not used by the priority, or low-latency queue, is not wasted; it is available to the bandwidth classes.

In IPsec transport mode, a G.729 voice call uses 48,000 bps (Layer 3 – 120 bytes * 50 pps * 8 = 48,000) versus 54,400 bps (Layer 3 – 136 bytes * 50 pps * 8 = 54,400). With one voice call active between the two handsets and VAD disabled, the following is an example **show interface** display output:

```
vpn18-2600-6#show interface se 0/1 | include rate
  Queueing strategy: weighted fair
  30 second input rate 50000 bits/sec, 50 packets/sec
  30 second output rate 50000 bits/sec, 50 packets/sec
```

In these configuration examples, the alternate or backup path is not used unless the primary path is unavailable. Both the logical path (the GRE tunnel) and the physical path are similar. Router *vpn18-2600-3* and *vpn18-2600-8* are the primary logical and physical path and *vpn18-2600-4* and *vpn18-2600-9* are the backup logical and physical path.

With this addressing scheme, recursive routing is addressed by more specific static routes targeted to the interface, while a supernet, 192.168.0.0/16 is advertised via EIGRP through the tunnel interface. Also note the core routers do not have a route to 192.168.6.1, the IPsec/GRE address for *vpn18-2600-6*, this is not an oversight, rather an illustration data traffic can be encrypted from network end to end without reachability to IPsec/GRE endpoints.

In the case of routers *vpn18-2600-8* and *vpn18-2600-9*, no QoS is enabled on these IPsec/GRE endpoints, QoS is addressed by the WAN aggregation routers *vpn18-2600-3* and *vpn18-2600-4*, as well as the remote branch router *vpn18-2600-2*. However, in the case of *vpn18-2600-10* and *vpn18-2600-6*, IPsec/GRE and QoS are configured on the same router. Either configuration is valid, however, from a design standpoint, separating QoS from IPsec/GRE head-end routers should be considered a more scalable and manageable approach.

Router Configuration—vpn18-2600-2

```

!
hostname vpn18-2600-2
!
boot system flash c2600-ik9s-mz.122-8.T
!
crypto isakmp key bigsecret address 192.168.8.1
crypto isakmp key bigsecret address 192.168.9.1
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
  mode transport
!
crypto map static-map local-address Loopback0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.8.1
  set transform-set vpn-test
  match address vpn-static1
crypto map static-map 20 ipsec-isakmp
  set peer 192.168.9.1
  set transform-set vpn-test
  match address vpn-static2
!
interface Loopback0
 ip address 192.168.2.1 255.255.255.252
!
interface Loopback1
 description Target address for dial peer
 ip address 10.0.3.1 255.255.255.0
!
interface Tunnel0
 ip address 10.0.100.1 255.255.255.0
 ip summary-address eigrp 44 10.0.0.0 255.255.252.0 5
 qos pre-classify
 tunnel source Loopback0
 tunnel destination 192.168.8.1 # Primary IPsec/GRE peer vpn18-2600-8
 crypto map static-map
!
interface Tunnel1
 ip address 10.0.101.1 255.255.255.0
 ip summary-address eigrp 44 10.0.0.0 255.255.252.0 5
 delay 60000 # Increasing the delay makes this the
 # backup peer
 qos pre-classify
 tunnel source Loopback0
 tunnel destination 192.168.9.1 # Backup IPsec/GRE peer vpn18-2600-9
 crypto map static-map
!
interface Serial0/0
 bandwidth 512
 no ip address
 encapsulation frame-relay
 frame-relay traffic-shaping

```

```

!
interface Serial0/0.100 point-to-point
 description Link to vpn18-2600-3
 bandwidth 512
 ip address 192.168.100.1 255.255.255.0
 frame-relay interface-dlci 100
 class ts-branch
 crypto map static-map
!
interface Serial0/0.101 point-to-point
 description Link to vpn18-2600-4
 bandwidth 512
 ip address 192.168.101.1 255.255.255.0
 frame-relay interface-dlci 101
 class ts-branch
 crypto map static-map
!
router eigrp 44
 network 10.0.0.0
 no auto-summary
 eigrp stub summary          # EIGRP stub configured
 eigrp log-neighbor-changes
!
! Two static routes to the head-end IPsec peers, 192.168.8.1 and
! 192.168.9.1 covered by the netmask of 255.255.254.0, the primary
! path to vpn18-2600-3 if available, otherwise use the second route
! with its higher administrative distance.
!
ip route 192.168.8.0 255.255.254.0 Serial0/0.100
ip route 192.168.8.0 255.255.254.0 Serial0/0.101 2
!
ip access-list extended vpn-static1
 permit gre host 192.168.2.1 host 192.168.8.1
ip access-list extended vpn-static2
 permit gre host 192.168.2.1 host 192.168.9.1
!
voice-port 1/0/0
 description 8-6666
!
dial-peer voice 10 voip
 destination-pattern 155467.....
 session target ipv4:10.251.0.1          # vpn18-2600-6
 ip qos dscp ef media
 ip qos dscp af31 signaling
 no vad
!
dial-peer voice 1 pots
 destination-pattern 15556786666
 port 1/0/0
!
end

```

Router Configuration—vpn18-2600-3

```

!
hostname vpn18-2600-3
!
boot system flash c2600-ik9s-mz.122-8.T
!
interface Serial0/0
 description link to vpn18-2600-4

```

```

bandwidth 2000
ip address 192.168.99.3 255.255.255.0
clockrate 2000000
!
interface FastEthernet0/1
description XTRA-NET
ip address 10.254.1.42 255.255.255.0
!
interface Serial0/1
bandwidth 512
no ip address
encapsulation frame-relay
frame-relay traffic-shaping
!
interface Serial0/1.100 point-to-point
description Link to vpn18-2600-2
bandwidth 512
ip address 192.168.100.2 255.255.255.0
frame-relay interface-dlci 100
class ts-headend
!
router eigrp 44
 redistribute static
 passive-interface Serial0/1.100
 network 10.0.0.0
 network 192.168.99.0
 network 192.168.100.0
 default-metric 64000 20000 255 1 1500
 no auto-summary
 eigrp log-neighbor-changes
!
! Create a /16 route to be advertised to vpn18-2600-6
!
ip route 192.168.0.0 255.255.0.0 Null0
!
! Redistribute the primary path into EIGRP, so vpn18-2600-4 will learn
! an EIGRP external dynamically.
!
ip route 192.168.2.0 255.255.255.0 Serial0/1.100
!
end

```

Router Configuration—vpn18-2600-4

```

!
hostname vpn18-2600-4
!
boot system flash c2600-ik9s-mz.122-8.T
!
interface Serial0/0
description link to vpn18-2600-3
bandwidth 2000
ip address 192.168.99.4 255.255.255.0
!
interface FastEthernet0/1
description XTRA-NET
ip address 10.254.1.46 255.255.255.0
!
interface Serial0/1
bandwidth 512
no ip address

```

```

encapsulation frame-relay
frame-relay traffic-shaping
!
interface Serial0/1.101 point-to-point
description link to vpn18-2600-2
bandwidth 512
ip address 192.168.101.2 255.255.255.0
frame-relay interface-dlci 101
class ts-headend
!
router eigrp 44
 redistribute static
 passive-interface Serial0/1.101
 network 10.0.0.0
 network 192.168.99.0
 network 192.168.101.0
 default-metric 64000 20000 255 1 1500
 no auto-summary
 eigrp log-neighbor-changes
!
!
! Create a /16 route to be advertised to vpn18-2600-6
!
ip route 192.168.0.0 255.255.0.0 Null0
!
! Due to admin distance of 240, this route will only be placed
! in the routing table if the EIGRP external (admin distance 170)
! from vpn18-2600-3 is withdrawn.
!
ip route 192.168.2.0 255.255.255.0 Serial0/1.101 240
!
end

```

Router Configuration—vpn18-2600-8

```

!
hostname vpn18-2600-8
!
boot system flash c2600-ik9s-mz.122-8.T
crypto isakmp key bigsecret address 192.168.2.1
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
 mode transport
!
crypto map static-map local-address Loopback0
crypto map static-map 10 ipsec-isakmp
 set peer 192.168.2.1
 set transform-set vpn-test
 match address vpn-static1
!
interface Loopback0
 ip address 192.168.8.1 255.255.255.0
!
interface Tunnel0
 ip address 10.0.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel destination 192.168.2.1 # vpn18-2600-2
 crypto map static-map
!
interface FastEthernet0/1
 description CORE

```

```

ip address 10.254.0.48 255.255.255.0
!
interface Ethernet1/0
  description XTRA-NET
  ip address 10.254.1.48 255.255.255.0
  crypto map static-map
!
router eigrp 44
  redistribute static
  network 10.0.0.0
  network 192.168.8.0
  default-metric 64000 20000 255 1 1500
  distribute-list 44 out Tunnel0
  no auto-summary
  eigrp log-neighbor-changes
!
! Create a /8 route to be advertised to the remote sites
!
ip route 10.0.0.0 255.0.0.0 Null0
!
ip access-list extended vpn-static1
  permit gre host 192.168.8.1 host 192.168.2.1
!
! Only need to send a /8 and /16 across the tunnel interface
!
access-list 44 permit 10.0.0.0
access-list 44 permit 192.168.0.0
access-list 44 deny any
!
end

```

Router Configuration—vpn18-2600-9

```

!
hostname vpn18-2600-9
!
boot system flash c2600-ik9s-mz.122-8.T
!
crypto isakmp key bigsecret address 192.168.2.1
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
  mode transport
!
crypto map static-map local-address Loopback0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.2.1
  set transform-set vpn-test
  match address vpn-static1
!
interface Loopback0
  ip address 192.168.9.1 255.255.255.0
!
interface Tunnel1
  ip address 10.0.101.2 255.255.255.0
  delay 60000
  tunnel source Loopback0
  tunnel destination 192.168.2.1 # vpn18-2600-2
  crypto map static-map
!
interface FastEthernet0/1
  description CORE

```

```

ip address 10.254.0.49 255.255.255.0
!
interface Ethernet1/0
  description XTRA-NET
  ip address 10.254.1.49 255.255.255.0
  crypto map static-map
!
router eigrp 44
  redistribute static
  network 10.0.0.0
  network 192.168.9.0
  default-metric 64000 20000 255 1 1500
  distribute-list 44 out Tunnel1
  no auto-summary
  eigrp log-neighbor-changes
!
! Create a /8 route to be advertised to the remote sites
!
ip route 10.0.0.0 255.0.0.0 Null0
!
ip access-list extended vpn-static1
  permit gre host 192.168.9.1 host 192.168.2.1
!
! Only need to send a /8 and /16 across the tunnel interface
!
access-list 44 permit 10.0.0.0
access-list 44 permit 192.168.0.0
access-list 44 deny any
!
end

```

Router Configuration—vpn18-2600-10

```

!
hostname vpn18-2600-10
!
boot system flash c2600-ik9s-mz.122-8.T
!
! Example of matching on DSCP rather than IP Precedence
!
class-map match-all call-setup
  description AF31
  match ip dscp 26
class-map match-any mission-critical
  description cs2 and cs6
  match ip dscp 16
  match ip dscp 48
class-map match-all voice
  description EF
  match ip dscp 46
!
policy-map hdlc
  class voice
    priority 672
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class class-default
    fair-queue
!

```

```

crypto isakmp key bigsecret address 192.168.6.1
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
  mode transport
!
crypto map HDLC local-address Loopback0
crypto map HDLC 10 ipsec-isakmp
  set peer 192.168.6.1
  set transform-set vpn-test
  match address hdlc-GRE
!
interface Loopback0
  ip address 192.168.10.1 255.255.255.0
!
interface Tunnel1
  ip address 10.249.0.2 255.255.255.0
  tunnel source Loopback0
  tunnel destination 192.168.6.1 # vpn18-2600-6
  crypto map HDLC
!
interface Serial0/0
  description to vpn-2600-6 se0/1
  bandwidth 2000
  ip address 192.168.65.2 255.255.255.0
  service-policy output hdlc
  clockrate 2000000
  crypto map HDLC
!
interface FastEthernet0/1
  description CORE
  ip address 10.254.0.50 255.255.255.0
!
router eigrp 44
  passive-interface Serial0/0
  network 10.0.0.0
  network 192.168.10.0
  distribute-list 44 out Tunnel1
  no auto-summary
  eigrp log-neighbor-changes
  no eigrp log-neighbor-warnings
!
! Need to know how to reach the crypto peer vpn18-2600-6
!
ip route 192.168.6.0 255.255.255.0 Serial0/0
!
ip access-list extended hdlc-GRE
  permit gre host 192.168.10.1 host 192.168.6.1
!
! Only need to send a /8 and /16 across the tunnel interface
!
access-list 44 permit 10.0.0.0
access-list 44 permit 192.168.0.0
access-list 44 deny any
!
end

```

Router Configuration—vpn18-2600-6

```

!
hostname vpn18-2600-6
!
boot system flash c2600-ik9s-mz.122-8.T
!
policy-map hdlc
  class voice
    priority 672
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class class-default
    fair-queue
!
crypto isakmp key bigsecret address 192.168.10.1
!
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
  mode transport
!
crypto map HDLC local-address Loopback0
crypto map HDLC 10 ipsec-isakmp
  set peer 192.168.10.1
  set transform-set vpn-test
  match address hdlc-GRE
!
interface Loopback0
  description target for IPSec
  ip address 192.168.6.1 255.255.255.0
!
interface Loopback1
  description target for VoIP dial-peers
  ip address 10.251.0.1 255.255.255.0
!
interface Tunnel1
  ip address 10.249.0.1 255.255.255.0
!
! Summarize up to the network core
!
ip summary-address eigrp 44 10.248.0.0 255.248.0.0 5
  qos pre-classify
  tunnel source Loopback0
  tunnel destination 192.168.10.1 # vpn18-2600-10
  crypto map HDLC
!
interface Serial0/1
  description to vpn-2600-10 se0/0
  bandwidth 2000
  ip address 192.168.65.1 255.255.255.0
  service-policy output hdlc
  crypto map HDLC
!
router eigrp 44
  passive-interface Serial0/1
  network 10.0.0.0
  no auto-summary
  eigrp stub summary
  eigrp log-neighbor-changes
  no eigrp log-neighbor-warnings
!
! Need to know how to reach crypto peer vpn18-2600-10

```



```
!  
ip route 192.168.10.0 255.255.255.0 Serial0/1  
!  
ip access-list extended hdlc-GRE  
  permit gre host 192.168.6.1 host 192.168.10.1  
voice-port 1/0/0  
  description 7-1111  
!  
dial-peer voice 10 voip  
  destination-pattern 155567.....  
  session target ipv4:10.0.3.1  
  ip qos dscp ef media  
  ip qos dscp af31 signaling  
  no vad  
!  
dial-peer voice 1 pots  
  destination-pattern 15546771111  
  port 1/0/0  
!  
end
```




Configuration Supplement—Dynamic Crypto Maps, Reverse Route Injection

This configuration supplement illustrates the head-end topology and configuration for an internal Cisco deployment of VoIP over IPSec for telecommuters in a SOHO implementation. The key features implemented in these examples include:

- Dynamic crypto maps (Digital Certificates required)
- IPSec appliances (router on a stick)
- Dual head-end routers for redundancy and availability
- Reverse Route Injection and IKE keepalive
- HSRP provides a next hop address for a firewall

An advantage of the use of dynamic crypto maps eliminates the need to make changes on the head-end routers as additional remote routers are deployed. The remote routers initiate the IPSec session to the peers defined in their configuration as devices on the remote LAN attempt to connect to resources at the corporate headquarters. In this implementation, Cisco 7960 IP Phones are installed on the SOHO LAN, their Skinny Protocol (connecting to TCP port 2000) with the headquarters CallManager generates *background* traffic, so the IKE and IPSec security associations are re-established at the end of their lifetimes. So the presence of the IP phone on the SOHO LAN allows the headquarters network administrator to telnet to the LAN interface of the SOHO router, even if the remote router is deployed in a home office and no one is home using the VPN connection.

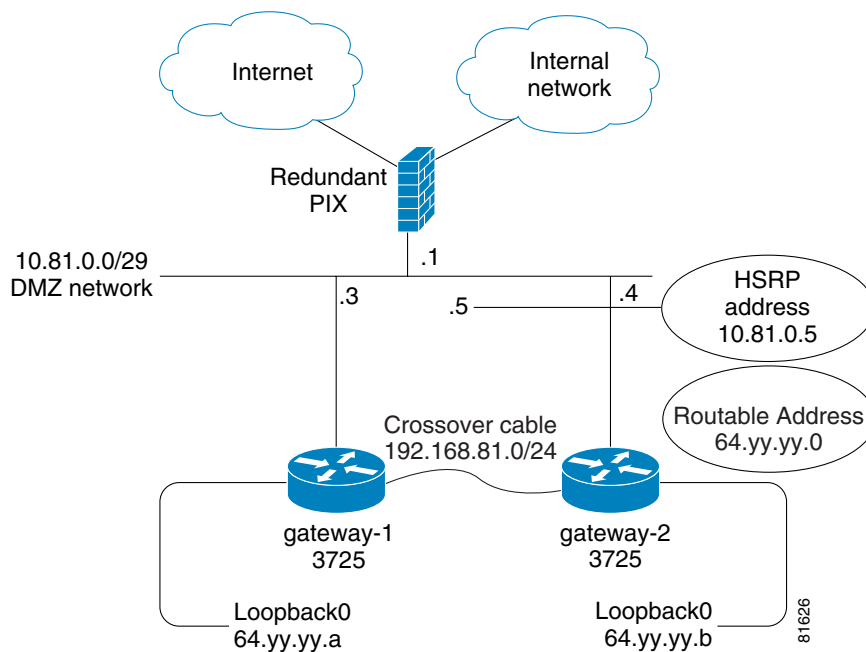
The head-end IPSec routers are configured and managed by an engineer in a separate organization from the IT support group which manages the Internet connectivity and firewall configuration. This is common in many enterprise networks, security responsibilities are routinely separated from internetworking support staff. As such, the head-end routers are truly *routers-on-a-stick*, as they have a single Fast Ethernet interface for connectivity to the enterprise core. The entire enterprise's traffic does not route through these head-end routers, only IP traffic which must be encrypted or decrypted between the headquarters and the remote SOHO users. Note the use of the **ip route-cache same-interface** command to avoid process switching in the router on a stick configuration.

The remote routers are configured with two IPSec peers defined in their crypto maps and IKE keepalives are configured to allow the remote routers a means to verify the availability of the head-end peers. During implementation it was observed that all the peers did not maintain an affinity to the first configured peer. The remote routers tend to *flip flop* between configured peers as short service disruptions to the home users cause IKE packet loss with the first configured peer. As the remote routers attempt to connect to the second peer, they often are successful with the second peer, so the IPSec session is established to that peer.

While the home user does not generally notice this *flip-flop*, it does present a routing issue if the IPsec session is established to the head-end router that is not the active HSRP router. To address this, the head-end routers are connected with a point to point Fast Ethernet interface and a simple CAT5 cross-over cable. Reverse Route Injection (RRI) is configured on the head-end crypto map and EIGRP is configured with the network address of the cross-over interfaces. Each head-end router redistributes the home user static routes inserted in the routing table by RRI. EIGRP advertises these injected routes to the other head-end router over the cross-over cable interface as an EIGRP external route (administrative distance 170).

The head-end topology is shown in [Figure C-1](#).

Figure C-1 Head-end Topology Diagram



Note

The IP addresses used in [Figure C-1](#) and configuration examples in this section (such as 64.yy.yy.a) *must* be public, routable IP addresses. They are shown here with alphabetic characters to avoid using real public addresses in the examples.

Visualize the path of a packet from the headquarters CallManager to the remote IP phone when the IPsec session is not on the active HSRP router. The firewall forwards the packet to the HSRP active router—IP address 10.81.0.5. This router, *gateway-1*, has an EIGRP external route in the routing table for the remote user's subnet, learned over the cross-over interface from *gateway-2*. The packet are forwarded to *gateway-2*, which is the active IPsec peer for that remote subnet. The routing table for *gateway-2* includes the RRI injected static route in the routing table. These static routes are recursive routes to the 0/0 (default) route out the Fast Ethernet interface (the *stick* interface) to the firewall. This interface has a crypto map configured. The packet matches the dynamic crypto map's access list and is encrypted in the IPsec tunnel to the remote router again via the firewall shown above.

Packets from the remote IP phone to the headquarters CallManager, are encrypted in the IPsec tunnel in our example which is peered with *gateway-2* public address 64.yy.yy.b. The firewall has a static route to this public address via 10.81.0.4. The packet is decrypted and routed out the same interface and back to the firewall to be routed to the CallManager in headquarters

If the above topology is implemented as shown, the only single point of failure is the cross-over cable between the two head-end IPsec routers. If the cable is removed or the interfaces fail, IPsec traffic to the standby HSRP router is *black holed*. Using two cross-over cables and two interfaces between the two head-end IPsec routers would eliminate the single point of failure.

Configure IKE keepalives on the head-end IPsec routers and the remote routers. Missed IKE keepalives on the head-end routers trigger removal of the RRI static routes

From the firewall perspective, only ESP (protocol 50) and IKE (UDP port 500) need be permitted to the loopback (public) interfaces of the two head-end IPsec routers, all other packets can be denied from the Internet. The Internet, Internal network and firewall topology is simplified for purposes of illustration.

The following configuration example shows the relevant routing and crypto configuration for each head-end router. The certificate portion of the crypto configuration is not shown for brevity.

```

!
hostname gateway-1
!
crypto isakmp policy 1
  encr 3des
crypto isakmp keepalive 10
!
crypto ipsec transform-set t1 esp-3des esp-sha-hmac
  mode transport
crypto ipsec transform-set t2 esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set t1 t2
reverse-route
!
crypto map test local-address Loopback0
crypto map test 1 ipsec-isakmp dynamic dmap
!
interface Loopback0
  description Public address
  ip address 64.yy.yy.a 255.255.255.255
!
interface FastEthernet0/0
  description Private
  ip address 10.81.0.3 255.255.255.248
  no ip redirects
ip route-cache same-interface
  standby 1 ip 10.81.0.5
  standby 1 priority 120
  standby 1 preempt
  crypto map test
!
interface FastEthernet0/1
  description X-Over cable to gateway-2
  ip address 192.168.81.3 255.255.255.0
!
router eigrp 64
redistribute static metric 1000 100 255 1 1500 route-map RRI
  passive-interface FastEthernet0/0
network 192.168.81.0
  no auto-summary
  eigrp log-neighbor-changes
  no eigrp log-neighbor-warnings
!
ip route 0.0.0.0 0.0.0.0 10.81.0.1
!
access-list 1 remark Home user address pool(s)

```

```

access-list 1 remark 10.81.2.0 / 23
access-list 1 remark 10.81.4.0 / 22
access-list 1 permit 10.81.2.0 0.0.1.255
access-list 1 permit 10.81.4.0 0.0.3.255
access-list 1 deny any
!
route-map RRI permit 10
  description Redistribute remote subnets from RRI
  match ip address 1
!
end

```

This sample configuration is for the second head-end router. When both routers are operational, this router is the standby HSRP router, as its HSRP priority is configured as 110 verses 120 for *gateway-1*.

```

!
hostname gateway-2
!
crypto isakmp policy 1
  encr 3des
crypto isakmp keepalive 10
!
crypto ipsec transform-set t1 esp-3des esp-sha-hmac
  mode transport
crypto ipsec transform-set t2 esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set t1 t2
  reverse-route
!
crypto map test local-address Loopback0
crypto map test 1 ipsec-isakmp dynamic dmap
!
interface Loopback0
  description Public address
  ip address 64.yy.yy.b 255.255.255.255
!
interface FastEthernet0/0
  description Private
  ip address 10.81.0.4 255.255.255.248
  no ip redirects
  ip route-cache same-interface
  standby 1 ip 10.81.0.5
  standby 1 priority 110
  standby 1 preempt
  crypto map test
!
interface FastEthernet0/1
  ip address 192.168.81.4 255.255.255.0
!
router eigrp 64
  redistribute static metric 1000 100 255 1 1500 route-map RRI
  passive-interface FastEthernet0/0
  network 192.168.81.0
  no auto-summary
  eigrp log-neighbor-changes
  no eigrp log-neighbor-warnings
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.81.0.1
!
access-list 1 remark Home user address pool(s)

```

```

access-list 1 remark 10.81.2.0 / 23
access-list 1 remark 10.81.4.0 / 22
access-list 1 permit 10.81.2.0 0.0.1.255
access-list 1 permit 10.81.4.0 0.0.3.255
access-list 1 deny any
!
route-map RRI permit 10
  description Redistribute remote subnets from RRI
  match ip address 1
!
end

```

Now an example of the relevant portion of a remote SOHO router. All packets from the LAN subnet 10.81.2.0/29 are encrypted. Note the two peers defined in the crypto map, these IP addresses are the Internet routable IP addresses from the two head-end IPsec peers. IKE connections are attempted in the order the peers appear in the crypto map.

```

!
hostname soho-vp
!
crypto isakmp policy 1
  encr 3des
crypto isakmp keepalive 10
!
crypto ipsec transform-set t1 esp-3des esp-sha-hmac
!
crypto map test 1 ipsec-isakmp
  description plain ol' ipsec
  set peer 64.yy.yy.a
  set peer 64.yy.yy.b
  set transform-set t1
  match address 101
  qos pre-classify
!
access-list 101 remark -----Crypto Map ACL-----
access-list 101 permit ip 10.81.2.0 0.0.0.7 any
!
end
!

```

Included here are displays from the remote and head-end routers to illustrate their configuration and normal state. First, the remote router. Note from the following crypto map display there are two peers defined and the current peer is indicated. In the IKE SA display, the destination (*dst*) address is the current peer's public (loopback 0) address and the source (*src*) address is the IP address on the outside Ethernet 0 interface. In this case the IP address is being assigned to this router by an upstream router that supports IPsec Pass-thru. This remote router is connected to the Internet via a 3rd party DSL modem.

```

router-vpn#show crypto map
Crypto Map "test" 1 ipsec-isakmp
  Description: plain ol' ipsec
  Peer = 64.yy.yy.a
  Peer = 64.yy.yy.b
  Extended IP access list 101
    access-list 101 permit ip 10.81.2.0 0.0.0.7 any
  Current peer: 64.yy.yy.a
  Security association lifetime: 4608000 kilobytes/3600 seconds
  PFS (Y/N): N
  Transform sets={
    t1,
  }
  QOS pre-classification
  Interfaces using crypto map test:
    Ethernet0

```

```
router-vpn#show crypto isakmp sa
dst          src          state          conn-id    slot
64.yy.yy.a   192.168.1.102  QM_IDLE       108        0
```

Displaying the IKE SA from the head-end peer router, notice the source address (*src*) 165.x.x.x this is the routable IP address of the above remote router, which 192.168.1.102 was translated to by the IPsec Pass-thru router.

```
gateway-1#show crypto isakmp sa
dst          src          state          conn-id    slot
64.yy.yy.a   165.x.x.x    QM_IDLE       16         0
```

From the head-end router, display the crypto map for the remote router. As this configuration uses dynamic crypto maps and the ISP assigns IP addresses for the remote routers, displaying the crypto map on the head-end router displays the access list that identifies packets to be encrypted. The access list entry identifies the remote subnet. For management and troubleshooting purposes, documenting the cross reference between remote subnet and the router's hostname is beneficial.

```
gateway-1#show crypto map
Crypto Map: "test" idb: Loopback0 local address: 64.yy.yy.a

Crypto Map "test" 1 ipsec-isakmp
  Dynamic map template tag: dmap
Crypto Map "test" 191 ipsec-isakmp
  Peer = 165.x.x.x
  Extended IP access list
    access-list permit ip any 10.81.2.0 0.0.0.7
    dynamic (created from dynamic map dmap/10)
  Current peer: 165.x.x.x
  Security association lifetime: 4608000 kilobytes/3600 seconds
  PFS (Y/N): N
  Transform sets={ t1, }
  Reverse Route Injection Enabled
```

Now look at the routing table. The RRI logic has inserted a static route for the remote subnet 10.81.2.0/29. There are two EIGRP external routes (10.81.2.24/29 and 10.81.2.48/29) learned from the second head-end router (*gateway-2*) as a result of the re-distribution of the RRI static routes. Looking at the routing table and the resulting static (or EIGRP external routes) provides a quick reference of the remote routers peered with this head-end router. Look for static routes for the remote subnets and the number of remote routers that are peered with the second head-end router (identified by the EIGRP external subnets).

```
gateway-1#show ip route

Gateway of last resort is 10.81.0.1 to network 0.0.0.0

    64.0.0.0/32 is subnetted, 1 subnets
C       64.yy.yy.a is directly connected, Loopback0
C       192.168.81.0/24 is directly connected, FastEthernet0/1
    10.0.0.0/29 is subnetted, 7 subnets
D EX   10.81.2.24 [170/2588160] via 192.168.81.4, 00:50:58, FastEthernet0/1
S       10.81.2.16 [1/0] via 0.0.0.0, FastEthernet0/0
S       10.81.2.0 [1/0] via 0.0.0.0, FastEthernet0/0
C       10.81.0.0 is directly connected, FastEthernet0/0
D EX   10.81.2.48 [170/2588160] via 192.168.81.4, 1d07h, FastEthernet0/1
S       10.81.2.40 [1/0] via 0.0.0.0, FastEthernet0/0
S       10.81.2.32 [1/0] via 0.0.0.0, FastEthernet0/0
S*    0.0.0.0/0 [1/0] via 10.81.0.1
```

Another perspective is found by comparing the RRI static routes on the second head-end route and the access-list entries from the crypto map.


```

gateway-2#show ip route

Gateway of last resort is 10.81.0.1 to network 0.0.0.0

    64.0.0.0/32 is subnetted, 1 subnets
C       64.yy.yy.b is directly connected, Loopback0
C       192.168.81.0/24 is directly connected, FastEthernet0/1
    10.0.0.0/29 is subnetted, 7 subnets
S       10.81.2.24 [1/0] via 0.0.0.0, FastEthernet0/0
D EX   10.81.2.16 [170/2588160] via 192.168.81.3, 1d08h, FastEthernet0/1
D EX   10.81.2.0 [170/2588160] via 192.168.81.3, 07:26:03, FastEthernet0/1
C       10.81.0.0 is directly connected, FastEthernet0/0
S       10.81.2.48 [1/0] via 0.0.0.0, FastEthernet0/0
D EX   10.81.2.40 [170/2588160] via 192.168.81.3, 05:48:36, FastEthernet0/1
D EX   10.81.2.32 [170/2588160] via 192.168.81.3, 03:51:48, FastEthernet0/1
S*    0.0.0.0/0 [1/0] via 10.81.0.1

gateway-2#show crypto map | include access-list
      access-list permit ip any 10.81.2.48 0.0.0.7
      access-list permit ip any 10.81.2.24 0.0.0.7

```

To test failover in the lab assessment, a remote user was in a phone conversation with an IP phone at the central site. The head-end router supporting the remote user (the standby HSRP router) was reloaded. It took approximately 23 seconds for the conversation to be audible in both directions. Then with the remote user peered with the HSRP active router, it was reloaded. It took approximately 31 seconds for the conversation to be audible in both directions. This test was conducted with only a few active remote routers. It was not a scale test, but rather a proof of concept test for this configuration.



A

anti-replay

- considerations [4-16](#)
- displaying drops [7-2](#)
- failures [3-6](#)

B

bandwidth

- allocation by link speed (table) [4-10](#)
- provisioning, traffic categorization (figure) [4-9](#)

bandwidth provisioning [4-5](#)

best practices [2-5](#)

branch office

- configuration listing
 - VPN router (Frame Relay) [A-5](#)
 - VPN router (HDLC) [A-8](#)

performance

- converged-traffic [5-11](#)
- product applicability by link speed [5-10](#)
- product selection [5-9](#)
- scalability [5-14](#)

C

CBWFQ

- Frame Relay considerations [4-8](#)

checklist

- design [4-35](#)
- implementation and configuration [6-24](#)

Cisco 7200VXR

- head-end [6-17](#)

Class Based Weighted Fair Queuing

See CBWFQ.

clear crypto isakmp command [7-10](#)

clear crypto sa command [7-10](#)

compression [3-5](#)

configuration

- applying service policy to HDLC T1 interfaces [6-16](#)
- applying traffic shaping [6-15](#)
- attaching service policy to Frame Relay map class [6-14](#)

branch office

- Frame Relay listing [A-5](#)

branch router

- HDLC listing [A-8](#)

checklist [6-24](#)

complete listings [A-1](#)

crypto map [6-21](#)

crypto map interface command [6-22](#)

DSCP [B-1](#)

EIGRP stub [B-1](#)

EIGRP summarization [6-2](#)

Frame Relay traffic shaping [6-11](#)

FRF.12 (LFI) [6-11](#)

GRE tunnels [6-2](#)

HDLC [B-1](#)

IKE [6-19](#)

IP GRE tunnel delay [6-3](#)

IPSec [6-19](#)

- local address [6-20](#)

- transform-set [6-21](#)

IPSec/IP GRE router [6-17](#)

ISAKMP [6-19](#)

ISAKMP policy [6-20](#)

network addressing [6-2](#)

pre-shared keys [6-20](#)

QoS [6-5](#)

- campus [6-5](#)
- class map [6-6](#)
- policy map [6-7](#)
- pre-classify [6-23](#)
- trust boundary [6-6](#)

scalability testbed files [A-1](#)

switching path [6-1](#)

tunnel interface [6-2, 6-3](#)

voice module listing [B-1](#)

VPN router listing [A-2](#)

WAN aggregation router [6-9](#)

crypto delay [3-5](#)

crypto engine

- current VoIP over IPsec capabilities [4-20](#)
- LLQ [4-21](#)
- QoS [4-20](#)

crypto map

- configuration [6-21](#)
- interface configuration [6-22](#)

crypto maps

- dynamic [C-1](#)

D

delay budget [3-5, 4-2](#)

design checklist [4-35](#)

DMVPN

- alternative to hub-and-spoke topology [3-6](#)

dynamic crypto maps [C-1](#)

Dynamic Multipoint VPN

See DMVPN.

E

E911/911

- handling calls [4-33](#)

EIGRP

- route summarization [6-2](#)
- verifying neighbors [7-3](#)

F

FIFO queue [3-6](#)

firewall considerations [4-16](#)

Frame Relay

- attaching to service policy [6-14](#)
- configuration [6-11](#)
- configuring traffic shaping [6-15](#)
- QoS considerations [4-8](#)
- traffic shaping [6-11](#)

Frame Relay Traffic Shaping

See FRTS.

FRF.12

- LFI implementation [4-8](#)

FRTS

- prerequisite for FRF.12 [4-8](#)

G

GRE

- IPsec tunnel considerations [4-14](#)
- IP tunnel configuration [6-2](#)
- tunnel delay configuration [6-3](#)

H

HDLC

- configuration supplement [B-1](#)
- service policy [6-16](#)
- T1 interfaces [6-16](#)

head-end

- availability and failover [5-6](#)
- Cisco 7200VXR [6-17](#)
- configuration listing

- VPN router [A-2](#)
- performance
 - converged-traffic [5-7](#)
- product selection [5-6](#)
- QoS effects [5-8](#)
- router locations [4-24](#)
- scalability [5-9](#)
- topology [4-23](#)

I

IKE

- clearing security associations [7-8, 7-10](#)
- configuration overview [6-19](#)

implementation

- checklist [6-24](#)
- considerations [6-1](#)
- summary [6-1](#)

Internet Key Exchange

See IKE

Internet Security Association and Key Management Protocol

See ISAKMP

ip load-sharing per-packet command [4-27](#)

IPSec

- anti-replay considerations [4-16](#)
- clearing security associations [7-8, 7-10](#)
- configuration overview [6-19](#)
- crypto engine QoS [4-20](#)
- firewall considerations [4-16](#)
- IP GRE router configuration [6-17](#)
- local address configuration [6-20](#)
- show command output [7-8](#)
- transform-set configuration [6-21](#)
- tunnel considerations [4-14](#)
- V3PN components [3-4](#)
- V3PN planning and design [4-14](#)

IP security

See IPSec

IP telephony

- calculating delay budget [4-2](#)
- hub-and-spoke [4-3](#)
- spoke-to-spoke [4-3](#)
- topology design options [4-3](#)
- V3PN components [3-1](#)
- V3PN planning and design [4-1](#)

ISAKMP

- policy configuration [6-20](#)

issues

- anti-replay failure [3-6](#)
- compressions [3-5](#)
- crypto delay [3-5](#)
- delay budget [3-5](#)
- FIFO queue [3-6](#)
- packet overhead [3-5](#)

L

Layer 3 packet size

- verifying [7-5](#)

LFI

- implemented by Layer 2 [4-8](#)

Link Fragmentation and Interleaving

See LFI.

link speed

- branch office products [5-10](#)

LLQ

- crypto engine [4-21](#)
- when required for crypto engine [4-22](#)

load sharing

- capabilities [4-27](#)
- design approach [4-28](#)
- head-end to branch [4-30](#)
- large flows [4-27](#)
- out-of-order packets [4-27](#)
- service provider considerations [4-32](#)
- summary [4-26](#)

low-latency queue

See LLQ.

M

mapping

ToS to CoS [6-5](#)

N

NetFlow

using to verify packet sizes [7-5](#)

using to verify ToS values [7-5, 7-6](#)

network address

configuration [6-2](#)

network topology

V3PN testbed diagram [A-1](#)

P

packet fragmentation [7-1](#)

packet overhead [3-5](#)

packet size

IPSec encrypted G.711 [4-7](#)

IPSec encrypted G.729 [4-5](#)

Layer 2 overhead [4-7](#)

performance

branch office

converged traffic [5-11](#)

head-end

converged traffic [5-7](#)

QoS effects [5-8](#)

policy map [6-7](#)

pre-shared keys

configuration [6-20](#)

product selection

branch office [5-9](#)

head-end [5-6](#)

overview [5-1](#)

performance and convergence requirements [5-15](#)

Q

QoS

affect on head-end [5-8](#)

bandwidth provisioning, WAN edge [4-5](#)

campus considerations [4-11](#)

configuration [6-5](#)

class map [6-6](#)

configuration [6-5](#)

policy map [6-7](#)

pre-classify [6-23](#)

trust boundary [6-6](#)

crypto engine [4-20](#)

Frame Relay considerations [4-8](#)

pre-classify [4-12](#)

show command output [7-12](#)

ToS byte [4-11](#)

ToS to CoS mapping [6-5](#)

traffic type allocation [4-9](#)

V3PN components [3-2](#)

V3PN planning and design [4-5](#)

Quality of Service

See QoS

R

recommendations

service provider [4-24](#)

requirements

convergence [5-15](#)

performance [5-15](#)

Reverse Route Injection. See RRI.

RRI [C-2, C-6](#)

S

scalability

- branch office [5-14](#)
- head-end [5-9](#)
- test methodology [5-2](#)

security

- clearing IPsec and IKE associations [7-10](#)
- crypto map configuration [6-21](#)
- IKE and IPsec configuration [6-19](#)
- IKE security associations [6-20](#)
- IPsec design [4-14](#)
- IPsec local address configuration [6-20](#)

service level agreement

See SLA.

service provider

- boundary consideration [4-24](#)
- cross-boundary considerations [4-25](#)
- service level agreements (SLA) [4-26](#)

show command

- IPsec output [7-8](#)
- QoS output [7-12](#)

show crypto engine connections active command [7-9](#)show crypto ipsec transform-set command [7-9](#)show crypto isakmp policy command [7-9](#)show crypto isakmp sa command [7-9](#)show crypto map command [7-8](#)show frame-relay fragment command [7-13](#)show policy-map interface command [7-12](#)

SLA

- minimum requirements

solution

- best practices guidelines [2-5](#)
- characteristics [2-4](#)
- overview [2-2](#)
- scope [1-4](#)
- V3PN components [3-1](#)

SRST

- V3PN design [4-33](#)

Survivable Remote Site Telephony

See SRST.

switching

- path configuration [6-1](#)

T

T1

- HDLC encapsulated interfaces [6-16](#)

ToS

- verifying values [7-6](#)

traffic profiles [5-3](#)

troubleshooting

- displaying anti-replay drops [7-2](#)
- IPsec show commands [7-8](#)
- packet fragmentation [7-1](#)
- verifying tunnel interfaces and EIGRP neighbors [7-3](#)

trust boundary [6-6](#)

tunnel interface

- configuring [6-2, 6-3](#)
- verifying [7-3](#)

V

V3PN

- anti-replay failure issue [3-6](#)
- best practices guidelines [2-5](#)
- compression issue [3-5](#)
- crypto delay issue [3-5](#)
- defined [1-1](#)
- delay budget issue [3-5](#)
- design
 - checklist [4-35](#)
 - scope [1-4](#)
- FIFO queue issue [3-6](#)
- head-end
 - router [4-24](#)
 - topology [4-23](#)

- implementation
 - issues [3-4](#)
- IPSec
 - components [3-4](#)
 - design [4-14](#)
- IP telephony
 - components [3-1](#)
 - design [4-1](#)
- key technologies [1-1](#)
- packet overhead issue [3-5](#)
- product selection [5-1](#)
- QoS
 - components [3-2](#)
 - design [4-5](#)
- site-to-site VPN [1-1](#)
- solution
 - characteristics [2-4](#)
 - components [3-1](#)
 - overview [2-2](#)
 - testbed diagram [A-1](#)
- traffic profiles [5-3](#)
- virtually-meshed VPN topology [3-6](#)

Voice and Video Enabled IPSec VPN

- See V3PN

W

- WAN
 - aggregation router configuration [6-9](#)
 - bandwidth provisioning [4-5](#)
 - Cisco 7200VXR router configuration [6-17](#)
 - implementation considerations [6-9](#)