



SECURITY

# IPSec VPN Design

The definitive design and deployment guide  
for secure virtual private networks

**Vijay Bollapragada, CCIE® No. 1606**  
**Mohamed Khalid, CCIE No. 2435**  
**Scott Wainner**



# IPSec VPN Design

**Vijay Bollapragada**  
**Mohamed Khalid**  
**Scott Wainner**

**Cisco Press**

800 East 96th Street  
Indianapolis, IN 46240 USA

## **IPSec VPN Design**

Vijay Bollapragada, Mohamed Khalid, Scott Wainner

Copyright© 2005 Cisco Systems, Inc.

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 2 3 4 5 6 7 8 9 0

Second Printing March 2006

Library of Congress Cataloging-in-Publication Number: 2002106378

ISBN: 1-58705-111-7

## **Warning and Disclaimer**

This book is designed to provide information about IPSec VPN design. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

## **Corporate and Government Sales**

Cisco Press offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales.

For more information, please contact U.S. Corporate and Government Sales, 1-800-382-3419, [corpsales@pearsontechgroup.com](mailto:corpsales@pearsontechgroup.com).

For sales outside the U.S., please contact International Sales at [international@pearsoned.com](mailto:international@pearsoned.com).

## **Trademark Acknowledgments**

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

## Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at [feedback@ciscopress.com](mailto:feedback@ciscopress.com). Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher	John Wait
Editor-in-Chief	John Kane
Cisco Representative	Anthony Wolfenden
Cisco Press Program Manager	Jeff Brady
Executive Editor	Brett Bartow
Production Manager	Patrick Kanouse
Development Editor	Grant Munroe
Project Editor	Sheila Schroeder
Copy Editor	Michelle Grandin
Technical Editors	Anthony Kwan, Suresh Subbarao, Michael Sullenberger
Team Coordinator	Tammi Barnett
Cover Designer	Louisa Adair
Composition	Mark Shirar
Indexer	Tim Wright



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

**European Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
[www-europe.cisco.com](http://www-europe.cisco.com)  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
[www.cisco.com](http://www.cisco.com)  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
Capital Tower  
168 Robinson Road  
#22-01 to #29-01  
Singapore 068912  
[www.cisco.com](http://www.cisco.com)  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic  
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy  
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal  
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden  
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Net Readiness Scorecard, Networking Academy, and ScriptShare are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCGE, 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, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, 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. (0303R)

Printed in the USA

## About the Authors

**Vijay Bollapragada**, CCIE No. 1606, is a director in the Network Systems Integration and Test Engineering group at Cisco Systems, where he works on the architecture, design, and validation of complex network solutions. An expert in router architecture and IP Routing, Vijay is a co-author of another Cisco Press publication titled *Inside Cisco IOS Software Architecture*. Vijay is also an adjunct professor in the Electrical Engineering department at Duke University.

**Mohamed Khalid**, CCIE No. 2435, is a technical leader working with IP VPN solutions at Cisco Systems. He works extensively with service providers across the globe and their associated Cisco account teams to determine technical and engineering requirements for various IP VPN architectures.

**Scott Wainner** is a Distinguished Systems Engineer in the U.S. Service Provider Sales Organization at Cisco Systems, where he focuses on VPN architecture and solution development. In this capacity, he works directly with customers in a consulting role by providing guidance on IP VPN architectures while interpreting customer requirements and driving internal development initiatives within Cisco Systems. Scott has more than 18 years of experience in the networking industry in various roles including network operations, network installation/provisioning, engineering, and product engineering. Most recently, he has focused his efforts on L2VPN and L3VPN service models using MPLS VPN, Pseudowire Emulation, and IPSec/SSL to provide VPN services to both enterprises and service providers. He holds a B.S. in Electrical Engineering from the United States Air Force Academy and a M.S. in Electronics and Computer Engineering from George Mason University in Fairfax, Virginia. Scott is currently an active member of the IEEE and the IETF.

---

## About the Technical Editors

**Anthony Kwan** is the director and executive project manager of infrastructure for HTA; CCNP, CCDP, MCSE, Master ASE, MCNE, CCIE(written). He has ten years of experience in the internetworking industry. He designed and built a number of secured enterprise datacenters with an upward budget of \$120 million. He also directed a number of consulting firms in building a Network Infrastructure and Technology consulting practice. He is a frequent contributor to Cisco Press and other publications specializing in networking technology. He can be reached at atonio888@yahoo.com.

**Suresh Subbarao** has worked in the networking area for the last 10 years. He is currently a network engineer at Cisco Systems focusing on security services for Service Providers with a special emphasis on IPsec VPNs.

**Michael Sullenberger** received a bachelor of science degree in mathematics from Harvey Mudd College in 1981. He started working with computer networks at the Stanford Linear Accelerator Center (SLAC) in 1981 as a Fortran programmer and as a user of the BITnet network, an early world wide 9600 baud network. At SLAC Michael also managed DEC VMS computers and gained knowledge of the DECnet and LAT protocol. He was also part of the introduction of Ethernet and FDDI networks to SLAC. In 1988 Michael moved to the networking group, where he assisted in transforming a large bridged, primarily DECnet, network to a routed multi-protocol, primarily TCP/IP, network. In 1994, he left SLAC to work for a small company, TGV, that wrote TCP/IP stacks and applications for OpenVMS and Windows systems. At TGV he worked in technical support where he learned the details of TCP/IP from the IP layer through the Application layer. TGV was bought by Cisco in 1996, and Michael moved into the Routing Protocols group, where he enhanced his knowledge of TCP/IP by adding information on the link-layer and IP routing protocols. In 1998, Michael moved to the Escalation Team at Cisco, where he continues to expand his TCP/IP knowledge in areas such as NAT, HSRP, GRE and IPsec Encryption. In 2000, he started a project, as the principle architect, that became the Cisco Dynamic Multipoint VPN (DMVPN) solution for scaling IPsec VPN networks. In 2004, the DMVPN solution won the Cisco Pioneer Award. Michael continues to this day working on enhancing DMVPN as well as designing and troubleshooting DMVPN and IPsec networks. Also starting in 2000 Michael has been a speaker each year at the Cisco Networkers Conferences in the area of site-to-site IPsec and DMVPN networks.

## Dedications

**Vijay Bollapragada:** To my best friend and wife, Leena, for her love and encouragement and for allowing me to take precious family time away to write this book. To my two lovely children, Amita and Abhishek, to my parents for instilling the right values in me, and all my wonderful friends.

Thanks to my coauthors, Mo and Scott, for bearing with me during the trials and tribulations of book writing and teaching me things along the way. And thanks to the awesome folks I work with at Cisco that constantly keep me challenged and remind me that there is something new to learn everyday.

**Mohamed Khalid:** First and foremost, I would like to acknowledge my parents—their dedication, sacrifice, and encouragement have been instrumental in all my achievements and success. Thanks to my wife Farhath, who gave me the time and constant encouragement to finish the book.

Thanks to Scott Wainner, Haseeb, and Sunil who provided valuable technical insights. Last but not least, I am deeply grateful to my friend and co-author, Vijay Bollapragada, who cajoled, encouraged, and assisted me in completing this book.

**Scott Wainner:** I would like to acknowledge my wife, Jill, for her love, patience, and encouragement. There are never enough hours in the day, so I thank her for caring for our family. I'd also like to thank my children—Craig, Brett, Natalie, and Caroline—for their patience and inspiration in exploring life's possibilities.

Special thanks go to my father and late mother—Tom and Zenith—for being an inspiration and guiding force in my life. To my colleagues, Vijay and Mo, you guys rock and it's been an honor working with you all these years. And finally, I'd like to acknowledge my God for granting me the gifts to fulfill this dream.

## Acknowledgments

This book would have not been possible without the help of many people whose many comments and suggestions improved the end result. First, we would like to thank the technical reviewers for the book, which include Anthony Kwan, Mike Sullenberger, and Suresh Subbarao. Their knowledge of the subject, attention to detail, and suggestions were invaluable. We would like to thank Brett Bartow of Cisco Press for constantly keeping the pressure and pulling all of this together. Without his help, this project would have never seen the light of day. We would also like to thank Grant Munroe and Chris Cleveland from Cisco Press for their attention to detail and editorial comments that improved the quality of the book tremendously. We would also like to thank the IPSec development team at Cisco—they are the ones that write and perfect the code that makes all the features discussed in this book possible.

*This page intentionally left blank*

## Contents at a Glance

Introduction xvi

<b>Chapter 1</b>	Introduction to VPNs	3
<b>Chapter 2</b>	IPSec Overview	11
<b>Chapter 3</b>	Enhanced IPSec Features	41
<b>Chapter 4</b>	IPSec Authentication and Authorization Models	89
<b>Chapter 5</b>	IPSec VPN Architectures	109
<b>Chapter 6</b>	Designing Fault-Tolerant IPSec VPNs	173
<b>Chapter 7</b>	Auto-Configuration Architectures for Site-to-Site IPSec VPNs	217
<b>Chapter 8</b>	IPSec and Application Interoperability	257
<b>Chapter 9</b>	Network-Based IPSec VPNs	293
<b>Index</b>		343

## Contents

	Introduction	xvi
<b>Chapter 1</b>	<b>Introduction to VPNs</b>	<b>3</b>
	Motivations for Deploying a VPN	3
	VPN Technologies	5
	Layer 2 VPNs	6
	Layer 3 VPNs	6
	GRE Tunnels	6
	MPLS VPNs	6
	IPSec VPNs	7
	Remote Access VPNs	8
	Summary	9
<b>Chapter 2</b>	<b>IPSec Overview</b>	<b>11</b>
	Encryption Terminology	11
	Symmetric Algorithms	12
	Asymmetric Algorithms	13
	Digital Signatures	14
	IPSec Security Protocols	15
	IPSec Transport Mode	16
	IPSec Tunnel Mode	17
	Encapsulating Security Header (ESP)	18
	Authentication Header (AH)	19
	Key Management and Security Associations	21
	The Diffie-Hellman Key Exchange	21
	Security Associations and IKE Operation	23
	IKE Phase 1 Operation	25
	Main Mode	26
	Aggressive Mode	27
	Authentication Methods	28
	IKE Phase 2 Operation	30
	Quick Mode	30
	IPSec Packet Processing	32
	Security Policy Database	32
	Security Association Database (SADB)	33
	Cisco IOS IPSec Packet Processing	34
	Summary	39
<b>Chapter 3</b>	<b>Enhanced IPSec Features</b>	<b>41</b>
	IKE Keepalives	41
	Dead Peer Detection	43
	Idle Timeout	47

---

Reverse Route Injection	50
RRI and HSRP	53
Stateful Failover	56
SADB Transfer	57
SADB Synchronization	57
IPSec and Fragmentation	65
IPSec and PMTUD	66
Look Ahead Fragmentation	69
GRE and IPSec	70
IPSec and NAT	76
Effect of NAT on AH	76
Effect of NAT on ESP	76
Effect of NAT on IKE	77
IPSec and NAT Solutions	77
NAT Traversal (NAT-T)	77
IPSec Pass-through	83
IKE Passing Through PAT	83
ESP Passing Through PAT	83
Restricted ESP Through PAT Mode	84
Summary	87
<b>Chapter 4</b> IPSec Authentication and Authorization Models	<b>89</b>
Extended Authentication (XAUTH) and Mode Configuration (MODE-CFG)	89
Mode-Configuration (MODECFG)	94
Easy VPN (EzVPN)	95
EzVPN Client Mode	96
Network Extension Mode	99
Digital Certificates for IPSec VPNs	103
Digital Certificates	103
Certificate Authority—Enrollment	104
Certificate Revocation	105
Summary	107
<b>Chapter 5</b> IPSec VPN Architectures	<b>109</b>
IPSec VPN Connection Models	109
IPSec Model	110
The GRE Model	111
The Remote Access Client Model	112
IPSec Connection Model Summary	112
Hub-and-Spoke Architecture	114
Using the IPSec Model	115
Transit Spoke-to-Spoke Connectivity Using IPSec	120

---

Internet Connectivity	126
Scalability Using the IPSec Connection Model	127
GRE Model	128
Transit Site-to-Site Connectivity	140
Transit Site-to-Site Connectivity with Internet Access	141
Scalability of GRE Hub-and-Spoke Models	143
Remote Access Client Connection Model	144
Easy VPN (EzVPN) Client Mode	145
EzVPN Network Extension Mode	151
Scalability of Client Connectivity Models	155
Full-Mesh Architectures	156
Native IPSec Connectivity Model	156
GRE Model	165
Summary	170
<b>Chapter 6</b> Designing Fault-Tolerant IPSec VPNs	<b>173</b>
Link Fault Tolerance	173
Backbone Network Fault Tolerance	174
Access Link Fault Tolerance	175
Multiple IKE Identities	176
Multiple IKE Identities Associated with Dial Backup	182
Single IKE Identity	183
Single IKE Identity Using Multi-link PPP on the Access Links	188
Access Link Fault Tolerance Summary	189
IPSec Peer Redundancy	189
Simple Peer Redundancy Model	189
Virtual IPSec Peer Redundancy Using HSRP	194
IPSec Stateful Failover	196
Peer Redundancy Using GRE	200
Virtual IPSec Peer Redundancy Using SLB	204
Server Load Balancing Concepts	205
IPSec Peer Redundancy Using SLB	205
Cisco VPN 3000 Clustering for Peer Redundancy	210
Peer Redundancy Summary	212
Intra-Chassis IPSec VPN Services Redundancy	212
Stateless IPSec Redundancy	213
Stateful IPSec Redundancy	213
Summary	214

**Chapter 7** Auto-Configuration Architectures for Site-to-Site IPSec VPNs 217

- IPSec Tunnel Endpoint Discovery 217
  - Principles of TED 218
  - Limitations with TED 220
  - TED Configuration and State 221
  - TED Fault Tolerance 225
- Dynamic Multipoint VPN 227
  - Multipoint GRE Interfaces 229
  - Next Hop Resolution Protocol 232
  - Dynamic IPSec Proxy Instantiation 236
  - Establishing a Dynamic Multipoint VPN 237
  - DMVPN Architectural Redundancy 248
  - DMVPN Model Summary 254
- Summary 255

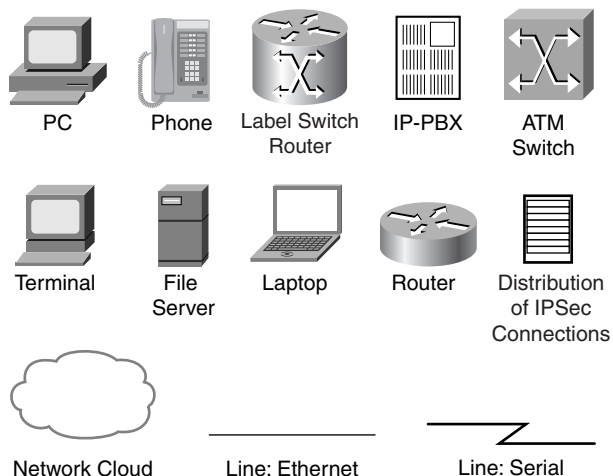
**Chapter 8** IPSec and Application Interoperability 257

- QoS-Enabled IPSec VPNs 258
  - Overview of IP QoS Mechanisms 258
  - IPSec Implications for Classification 259
    - QoS Applied to IPSec Transport Mode 260
    - QoS Applied to IPSec Tunnel Mode 261
    - IPSec Transport Mode - QoS Attribute Preservation of GRE Tunnels 261
    - Transitive QoS Applied to IPSec 264
    - Internal Preservation of QoS Attributes 264
  - IPSec Implications on QoS Policies 266
    - IPSec Implications of Packet Size Distribution on Queue Structures 266
    - IPSec Implications of Packet Size on Queue Bandwidth Assignments 266
- VoIP Application Requirements for IPSec VPN Networks 267
  - Delay Implications 267
  - Jitter Implications 269
  - Loss Implications 270
    - Mitigating Anti-replay Loss in Combined Voice/Data Flows 270
    - Mitigating Anti-replay Loss in Separate Voice/Data Flows 270
    - Engineering Best Practices for Voice and IPSec 271
- IPSec VPN Architectural Considerations for VoIP 271
  - Decoupled VoIP and Data Architectures 272
  - VoIP over IPSec Remote Access 274
  - VoIP over IPSec-Protected GRE Architectures 275
  - VoIP Hub-and-Spoke Architecture 277
  - VoIP over DMVPN Architecture 278
    - VoIP Bearer Path Optimization with DMVPN 279
    - VoIP Bearer Path Synchronization with DMVPN 279
  - VoIP Traffic Engineering Summary 279

---

Multicast over IPSec VPNs	280
Multicast over IPSec-protected GRE	280
Multicast on Full-Mesh Point-to-Point GRE/IPSec Tunnels	282
DMVPN and Multicast	285
Multicast Group Security	287
Group Security Key Management	287
Group Security Association	289
Multicast Group Security Summary	291
Multicast Encryption Summary	291
Summary	291
<b>Chapter 9</b> Network-Based IPSec VPNs	<b>293</b>
Fundamentals of Network-Based VPNs	293
The Network-Based IPSec Solution: IOS Features	296
The Virtual Routing and Forwarding Table	296
Crypto Keyrings	297
ISAKMP Profiles	297
Operation of Network-Based IPSec VPNs	299
A Single IP Address on the PE	300
Front-Door and Inside VRF	300
Configuration and Packet Flow	301
Generic MPLS VPN Configuration on the PE	305
Mapping an IPSec Tunnel from a Site into IVRF at the PE	306
Mapping an IPSec Tunnel from a Telecommuter into an IVRF at the PE	315
Termination of IPSec on a Unique IP Address Per VRF	321
Network-Based VPN Deployment Scenarios	324
IPSec to MPLS VPN over GRE	324
DMVPN and VRF	327
IPSec to L2 VPNs	330
PE-PE Encryption	334
Summary	339
<b>Index</b>	<b>343</b>

## Icons Used in This Book



## Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italics* indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [ ] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

## Introduction

VPNs are becoming more important for both enterprises and service providers. IPsec specifically is one of the more popular technologies for deploying IP-based VPNs. There are many books in the market that go into technical details of IPsec protocols and cover product level configuration, but they do not address overall design issues for deploying IPsec VPNs.

## The Goals of This Book

The objective of this book is to provide you with a good understanding of design and architectural issues of IPsec VPNs. This book will also give you guidance on enabling value-added services and integrating IPsec VPNs with other Layer 3 (MPLS VPN) technologies.

## Who Should Read This Book

The primary audience for this book is network engineers involved in design, deployment, and troubleshooting of IPsec VPNs. The assumption in this book is that you have a good understanding of basic IP routing, although IPsec knowledge is not a prerequisite.

## How This Book Is Organized

The book is divided into three general parts. Part I covers the general architecture of IPsec, including its protocols and Cisco IOS IPsec implementation details. Part II, beginning with Chapter 5, examines the IPsec VPN design principles covering hub-and-spoke, full-mesh, and fault-tolerant designs. Part II also covers dynamic configuration models used to simplify IPsec VPNs designs, and presents a case study. Part III, beginning with Chapter 8, covers design issues in adding services to an IPsec VPN such as voice, multicast, and integrating IPsec VPNs with MPLS VPNs. The book is organized as follows:

- **Part I, “Introduction and Concepts”**
  - **Chapter 1, “Introduction to VPNs”**—Provides an introduction to VPN concepts and covers a brief introduction to various VPN technologies.
  - **Chapter 2, “IPsec Overview”**—Gives an overview of IPsec protocols and describes differences between transport mode and tunnel mode. Cisco IOS IPsec packet processing is also explained in this chapter.
  - **Chapter 3, “Enhanced IPsec Features”**—Introduces advanced IPsec features that improve IPsec VPN scalability and fault tolerance, such as dead peer detection and control plane keepalives. This chapter also explains the challenges of IPsec interoperating with Network Address Translation (NAT) and Path Maximum Transmission Unit detection (PMTUD) and how to overcome these challenges.

- 
- **Chapter 4, “IPSec Authentication and Authorization Models”**—Explores IPSec features that are primarily called upon for the remote access users such as Extended Authentication (XAUTH) and Mode-configuration (MODE-CFG). It also explains the Cisco EzVPN connection model and digital certificate concepts.
  - **Part II, “Design and Deployment”**
    - **Chapter 5, “IPSec VPN Architectures”**—Covers various IPSec connections models such as native IPSec, GRE, and remote access. Deployment architectures for each of the connection models are explored with pros and cons for each architecture.
    - **Chapter 6, “Designing Fault-Tolerant IPSec VPNs”**—Discusses how to introduce fault tolerance into VPN architectures and describes the caveats with the various fault-tolerance methods.
    - **Chapter 7, “Auto-Configuration Architectures for Site-to-Site IPSec VPNs”**—Covers mechanisms to alleviate the configuration complexity of a large-scale IPSec VPN; Tunnel Endpoint Discovery (TED) and Dynamic Multipoint VPNs (DMVPN) are the two mechanisms discussed in depth.
  - **Part III, “Service Enhancements”**
    - **Chapter 8, “IPSec and Application Interoperability”**— Examines the issues with IPSec VPNs in the context of the running applications such as voice and multicast over the VPN.
    - **Chapter 9, “Network-Based IPSec VPNs”**—Concludes by introducing the concept of network-based VPNs.





# Introduction to VPNs

---

Virtual private networks, commonly referred to as VPNs, are not an entirely new concept in networking. As the name suggests, a VPN can be defined as a private network service delivered over a public network infrastructure. A telephone call between two parties is the simplest example of a virtual private connection over a public telephone network. Two important characteristics of a VPN are that it is virtual and private.

There are many types of VPNs, such as Frame Relay and ATM, and entire books can and have been written about each of these VPN technologies. The focus of this book is on a VPN technology known as *IPSec*.

## Motivations for Deploying a VPN


This chapter introduces some of the VPN technologies and helps to explain the motivations for deploying a VPN. The primary reason for deploying a VPN is cost savings.

Corporations with offices all over the world often need to interconnect them in order to conduct everyday business. For these connections, they can either use dedicated leased lines that run between the offices or have each site connect locally to a public network, such as the Internet, and form a VPN over the public network.

Figure 1-1 shows an international corporation that connects to each site using leased lines. Each connection is point-to-point and requires a dedicated leased line to connect it to another site. If each site needs to be connected to every other site (a situation also known as any-to-any or full-mesh connectivity),  $n-1$  leased lines would be required at each site where  $n$  is the number of sites. Leased lines are typically priced based on the distance between the sites and bandwidth offered. Cross-country and intercontinental links are typically very expensive, making full-mesh connectivity with leased lines very expensive.

Figure 1-2 shows an alternate method of connecting the same sites of the corporation, this time over a public network such as the Internet. In this model, each site is connected to the public network at its closest point, possibly via a leased line, but all connections between sites are virtual connections. The cloud in the figure represents a virtual connection between the sites, as opposed to a physical dedicated connection between sites in the leased-line model.

**Figure 1-1** *Connecting Sites of a Corporation over Leased Lines*



---

**NOTE**


A public network can be defined as a network with an infrastructure shared by many users of that network. Bear in mind that the word “public” does not mean that the network is available free to anyone. Many service providers have large ATM and Frame Relay public networks, and the Internet is probably the most ubiquitous public network of them all.

---

Although connecting the sites over a public network has obvious cost advantages over the dedicated leased line model and provides significant cost savings to the corporation, this model also introduces risks, such as the following:

- Data security
- Lack of dedicated bandwidth between sites

In the VPN model, the corporation’s data is being transported across a public network, which means other users of the public network can potentially access the corporation’s data and thereby pose a security risk.

**Figure 1-2** *Connecting Sites of a Corporation over a Public Network*

The second risk in the VPN model is the lack of dedicated bandwidth availability between sites that the leased line model provides. Because the VPN model connects sites using a virtual connection and the physical links in the public network are shared by many sites of many different VPNs. Bandwidth between the sites is not guaranteed unless the VPN allows some form of connection admission control and bandwidth reservation schemes. Both risks can be mitigated—the next section introduces some VPN technologies that overcome these risks.

## VPN Technologies

In the simplest sense, a VPN connects two endpoints over a public network to form a logical connection. The logical connections can be made at either Layer 2 or Layer 3 of the OSI model, and VPN technologies can be classified broadly on these logical connection models as Layer 2 VPNs or Layer 3 VPNs. Conceptually, establishing connectivity between sites over a Layer 2 or Layer 3 VPN is the same. The concept involves adding a “delivery header” in front of the payload to get it to the destination site. In the case of Layer 2 VPNs, the delivery header is at Layer 2, and in the case of Layer 3 VPNs, it is (obviously) at Layer 3. ATM and Frame Relay are examples of Layer 2 VPNs; GRE, L2TP, MPLS, and IPSec are examples of Layer 3 VPN technologies.

## Layer 2 VPNs

Layer 2 VPNs operate at Layer 2 of the OSI reference model; they are point-to-point and establish connectivity between sites over a virtual circuit. A virtual circuit is a logical end-to-end connection between two endpoints in a network, and can span multiple elements and multiple physical segments of a network. The virtual circuit is configured end-to-end and is usually called a permanent virtual circuit (PVC). A dynamic point-to-point virtual circuit is also possible and is known as a switched virtual circuit (SVC); SVCs are used less frequently because of the complexity involved in troubleshooting them. ATM and Frame Relay are two of the most popular Layer 2 VPN technologies. ATM and Frame Relay providers can offer private site-to-site connectivity to a corporation by configuring permanent virtual circuits across a shared backbone.

One of the advantages of a Layer 2 VPN is the independence of the Layer 3 traffic payload that can be carried over it. A Frame Relay or ATM PVC between sites can carry many different types of Layer 3 traffic such as IP, IPX, AppleTalk, IP multicast, and so on. ATM and Frame Relay also provide good quality of service (QoS) characteristics, which is especially critical for delay-sensitive traffic such as voice.

## Layer 3 VPNs

A connection between sites can be defined as a Layer 3 VPN if the delivery header is at Layer 3 of the OSI model. Common examples of Layer 3 VPNs are GRE, MPLS, and IPSec VPNs. Layer 3 VPNs can be point-to-point to connect two sites such as GRE and IPSec, or may establish any-to-any connectivity to many sites using MPLS VPNs.

### GRE Tunnels

Generic routing encapsulation (GRE) was originally developed by Cisco and later standardized as RFC 1701. An IP delivery header for GRE is defined in RFC 1702. A GRE tunnel between two sites that have IP reachability can be described as a VPN, because the private data between the sites is encapsulated in a GRE delivery header.

Because the public Internet is probably the most ubiquitous public network in the world, it is possible to connect many sites of a corporation using GRE tunnels. In this model, each site of the corporation requires only physical connectivity to its Internet service provider, as all of the connections between sites are over GRE tunnels. Although VPNs built over the Internet using GRE are possible, they are rarely used for corporate data due to the inherent risks and lack of strong security mechanisms associated with GRE.

### MPLS VPNs

Pioneered by Cisco, Multiprotocol Label Switching was originally known as Tag Switching and later standardized via the IETF as MPLS. Service providers are increasingly deploying MPLS

to offer MPLS VPN services to customers. A common principle among all VPN technologies is encapsulation of private data with a delivery header; MPLS VPNs use labels to encapsulate the original data, or payload, to form a VPN between sites.

---

**NOTE**

Creating an MPLS VPN is the most popular application and the primary motivation for deploying MPLS; other applications of MPLS include traffic engineering offering Layer 2 VPN services over MPLS.

---

RFC 2547 defines a scheme for offering VPN service using MPLS. One of the key advantages of MPLS VPNs over other VPN technologies is that it offers the flexibility to configure arbitrary topologies between VPN sites. For example, if three sites of a corporation all must be connected to one another in a full-mesh (any-to-any) configuration using ATM, Frame Relay, GRE, or IPSec technologies, each site requires two virtual circuits, or tunnels, to every other site. The addition of a fourth site to this full-mesh configuration requires that three tunnels, or virtual circuits, exist at each site, and calls for modification in the configurations at all the sites. If  $n$  is the number of sites in a VPN, the configuration complexity for this model is  $O(n)$  and the scalability is  $O(n^2)$ . If the same three sites are connected over an MPLS VPN, the addition of the fourth site requires configuration change at only the fourth site. The configuration complexity of this model with  $n$  sites is always a constant and is  $O(1)$ .

The fact that there are virtually no point-to-point tunnels for connecting sites of an MPLS VPN renders them very scalable. Any-to-any connectivity between sites of a VPN and extranet connectivity across VPNs are easy to achieve using MPLS VPNs compared to other tunneling techniques such as GRE. One of the drawbacks of an MPLS VPN is that connectivity between the sites of a VPN is restricted to sites where the provider has points of presence. Although a GRE tunnel could be used across the Internet to extend its reach, GRE by itself has minimal security. We address this issue in Chapter 9, “Network-Based IPSec VPNs.”

## IPSec VPNs

One of the main concerns for anyone using any VPN is security of data when it traverses a public network. In other words, how does one prevent malicious eavesdropping of data in a VPN?

Encrypting the data is one way to protect it. Data encryption may be achieved by deploying encryption/decryption devices at each site. IPSec is a suite of protocols developed under the auspices of the IETF to achieve secure services over IP packet-switched networks. The Internet is the most ubiquitous packet-switched public network; therefore, an IPSec VPN deployed over the public Internet can mean significant cost savings to a corporation as compared to a leased-line VPN.

IPSec services allow for authentication, integrity, access control, and confidentiality. With IPSec, the information exchanged between remote sites can be encrypted and verified. Both remote access clients and site-to-site VPNs can be deployed using IPSec. Subsequent chapters focus on the IPSec protocols and deployment models that use IPSec.

## Remote Access VPNs

As stated previously, VPNs can be classified into site-to-site VPNs and remote access VPNs. Frame Relay, ATM, GRE, and MPLS VPN can be considered site-to-site VPNs because information relevant to the configuration between sites is known in advance at both sides and, more importantly, are static and therefore do not change dynamically. On the other hand, consider a telecommuter who needs VPN access to corporate data over the Internet. The information required to establish a VPN connection such as an IP address of the telecommuter changes dynamically depending on the location of the telecommuter and is not known in advance to the other side of the VPN. This type of VPN can be classified as a remote access VPN.

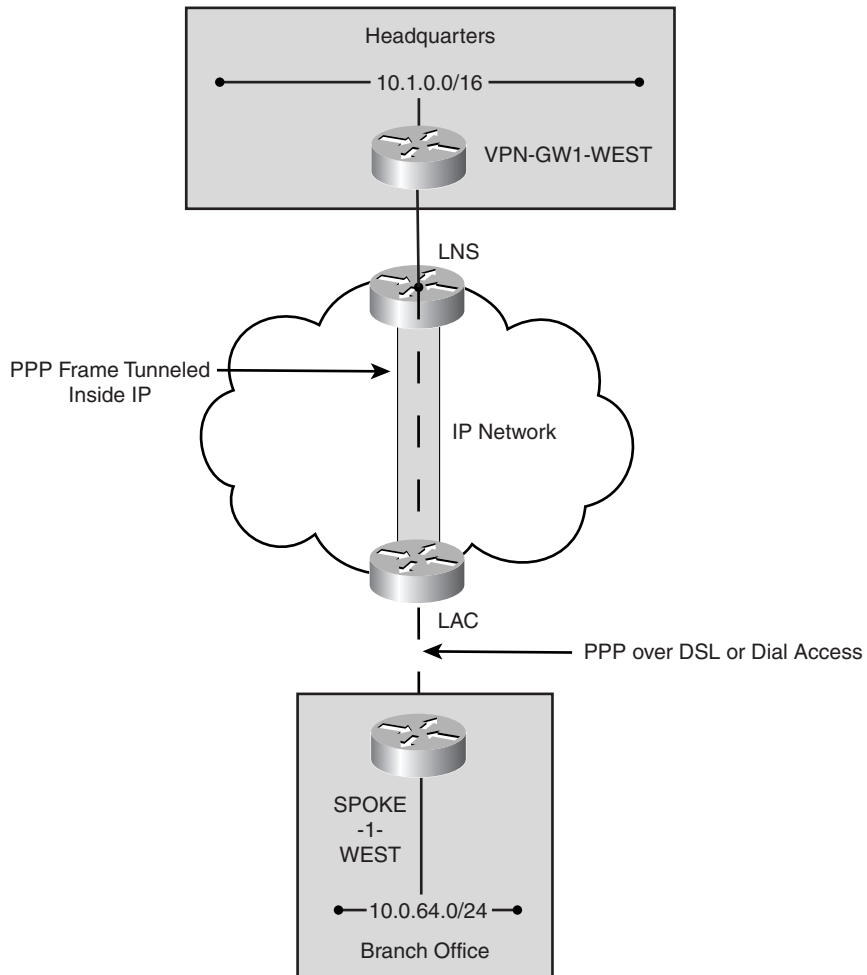
Remote access to corporate data resources has been a critical enabler for improved productivity, especially for mobile workers. Telecommuters, “road warriors,” and remote offices rely on timely access to mission-critical information in order to maintain a competitive advantage in the marketplace. The reliance on remote access has driven demand for higher capacity connections with extended durations from end users. As a result, increased costs are incurred, primarily in the form of telephony charges, for access to the corporate data.

Although dial-up networking provides a universal local access solution, it can be very expensive for long distance and metered local access calls. Remote access VPN connections provide the best solution, mitigating metered telephone charges while allowing the corporation to leverage new last-mile access technologies such as cable and DSL.

Two of the most common remote access methods for VPN access are Layer 2 tunneling protocol (L2TP) and IPSec. L2TP is an IETF standard (RFC 2661) for transporting PPP frames over IP. L2TP provides dial-up users with a virtual connection to a corporate gateway over an IP network, which could be the Internet. Figure 1-3 shows the L2TP model.

The remote user initiates a PPP session to the closest access server, known as a local access concentrator (LAC) via a local telephone call. The LAC authenticates the remote user and determines which local network server (LNS) will terminate the remote user. An L2TP tunnel is established between the LAC and the LNS, and once the LNS authenticates the user, a virtual interface for PPP termination is created on the LNS analogous to a direct-dialed connection to the LNS.

IPSec is another VPN technology that can be used to connect remote access users. This entire book is devoted to the topic of IPSec VPNs, and remote access is specifically covered in detail in Chapter 4, “IPSec Authentication and Authorization Models.”

**Figure 1-3** *Remote Access VPN Using L2TP*

## Summary

In this brief introduction to VPNs, you learned that network designers can choose from a wide range of technologies to create VPNs which can be classified into Layer 2 or Layer 3 VPNs, and further into site-to-site and remote access VPNs. Technologies such as Frame Relay, ATM, GRE, and MPLS are used with site-to-site VPNs. The most common remote access VPN technology is L2TP, and IPsec can be used for both remote access and site-to-site VPNs.



# IPSec Overview

---

Chapter 1, “Introduction to VPNs,” introduced VPN concepts at a high level and presented an overview of several technologies that use VPNs. In this chapter, you will explore the building blocks of an IPSec VPN and obtain an understanding of IPSec architecture and how the various components of IPSec interact with each other to create a VPN. You will also look at some Cisco-specific IPSec implementation details and how IPSec packet processing is performed on Cisco IOS platforms.

A common misconception about IPSec is that it is a single protocol for providing these security services for IP traffic. In fact, IPSec is really a *suite*, or collection, of protocols for security defined by the IPSec working group in the IETF. The baseline IPSec architecture and fundamental components of IPSec are defined in RFC2401 as the following:

- **Security protocols**— Authentication header (AH) and encapsulation security payload (ESP)
- **Key management**—ISAKMP, IKE, SKEME
- **Algorithms**—for encryption and authentication

The interaction between these components of IPSec is intertwined in such a way that it is a bit hard to understand one of the components without understanding another. A quote from a draft submitted to the IPSec IETF working group sums it up pretty well: “Perhaps IPSec is well understood by some, but frequent questions on the developers’ mailing list confirm that one cannot become an IPSec expert merely by reading the RFCs. Much valuable information is buried deep in the list archives or in the minds of its designers.”

You will start your IPSec journey with an introduction to encryption terminology, followed by an examination of the IPSec security protocols (AH and ESP), and lastly, an explanation of security associations and key management.

## Encryption Terminology

Security and data confidentiality are prime requirements for any VPN. One of the primary reasons for choosing IPSec as your VPN technology is the confidentiality of data provided by the encryption that is built in.

**NOTE**

*Encryption* is the transformation of plain text into a form that makes the original text incomprehensible to an unauthorized recipient that does not hold a matching key to decode or decrypt the encrypted message.

*Decryption* is the reverse of encryption; it is the transformation of encrypted data back into plain text. Encryption techniques are as old as history—in fact, Julius Cæsar apparently did not trust his messengers and therefore encrypted his military messages to his generals with a simple encryption scheme; he replaced every A by D, every B by E, and so on. Only someone who knew the *key* (to shift each alphabetical letter by three, in this case) would be able to decrypt the message.

---

A *cryptographic algorithm*, also called a *cipher*, is the mathematical function used for encryption and decryption. Generally, there are two related functions—one for encryption and the other for decryption. Security of data in modern cryptographic algorithms is based on the *key* (or keys). It doesn't matter if an eavesdropper knows your algorithm; if he or she doesn't know your particular key, an eavesdropper will be unable read your messages.

Cryptographic algorithms can be classified into two categories:

- Symmetric
- Asymmetric

## Symmetric Algorithms

Symmetric cryptographic algorithms are based on the sender and receiver of the message knowing and using the same secret key. The sender uses a secret key to encrypt the message, and the receiver uses the same key to decrypt it. The main problem with using the symmetric key approach is finding a way to distribute the key without anyone else finding it out. Anyone who overhears or intercepts the key in transit can later read and modify messages encrypted or authenticated using that key, and can forge new messages. DES, 3DES, and AES are popular symmetric encryption algorithms. A detailed explanation of these algorithms is beyond the scope of this book.

**NOTE**

DES uses a 56-bit key and is not considered secure anymore; in 1999, the DES key was cracked in less than 24 hours by using an exhaustive key search. Triple DES (3DES) and AES are the recommended encryption algorithms as of this writing.

---

## Asymmetric Algorithms

Asymmetrical encryption algorithms, also known as public key algorithms, use separate keys—one for encryption and another for decryption. The encryption key is called the *public key* and can be made public. Only the *private key*, used for decryption, needs to be kept secret. Although the public and private keys are mathematically related, it is not feasible to derive one from the other. Anyone with a recipient’s public key can encrypt a message, but the message can only be decrypted with a private key that only the recipient knows. Therefore, a secure communication channel to transmit the secret key is no longer required as in the case of symmetric algorithms.

**Figure 2-1** *Public Key Encryption*

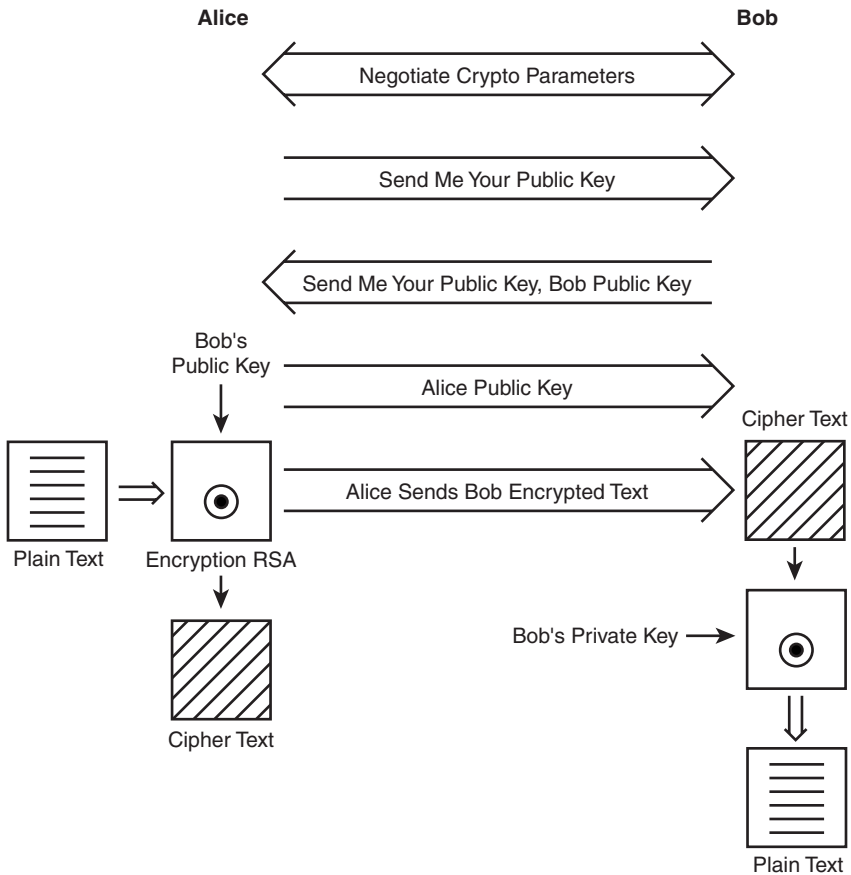


Figure 2-1 illustrates how public key encryption algorithms work. Bob and Alice communicate securely using public key encryption as follows:

- 1 Alice and Bob agree on a public key algorithm.
- 2 Bob sends Alice his public key and Alice sends Bob her public key.
- 3 Alice sends Bob a message, encrypting the message using Bob's public key.
- 4 Bob receives the message and decrypts Alice's message using his private key.

---

**NOTE**

Whenever an encryption theory or algorithm is used to describe a transaction between two parties, longstanding tradition has it that the parties are called Alice and Bob, and the eavesdropper in the middle is called Eve or Blackhat. Rumor has it that early on, the FBI and CIA actually went looking for Alice and Bob, because they were passing so many encrypted messages.

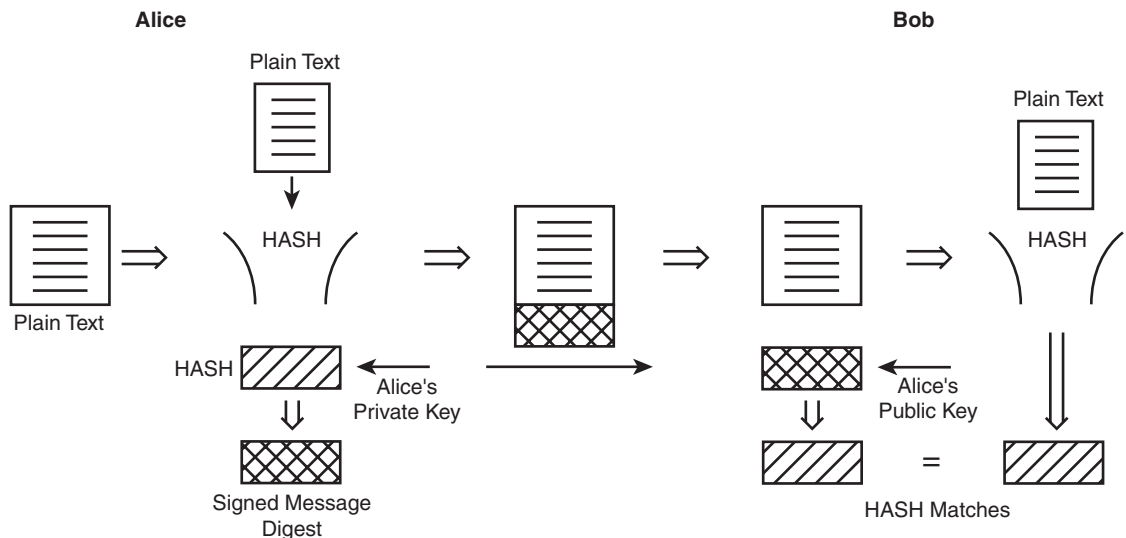
---

In reality, public key encryption is rarely used to encrypt messages because it is much slower than symmetric encryption; however, public key encryption is used to solve the problem of key distribution for symmetric key algorithms, which is, in turn used to encrypt actual messages. Therefore, public key encryption is not meant to replace symmetric encryption, but can supplement it and make it more secure.

## Digital Signatures

Another good use of public key encryption is for message authentication, also known as a digital signature.

Encrypting a message with a private key creates a digital signature, which is an electronic means of authentication and provides non-repudiation. *Non-repudiation* means that the sender will not be able to deny that he or she sent the message. That is, a digital signature attests not only to the contents of a message, but also to the identity of the sender. Because it is usually inefficient to encrypt an actual message for authentication, a document *hash* known as a message digest is used. The basic idea behind a message digest is to take a variable length message and convert it into a fixed length compressed output called the message digest. Because the original message cannot be reconstructed from the message digest, the hash is labeled "one-way." Alice and Bob's communication using digital signature is shown in Figure 2-2.

**Figure 2-2** *Signed Message Digest*

- 1 Alice computes a one-way hash of a document that she wishes to send Bob.
- 2 Alice encrypts the hash with her private key. The encrypted message digest becomes the digital signature.
- 3 Alice sends the document along with the digital signature to Bob.
- 4 Bob decrypts the digital signature using Alice's public key and also computes a one-way hash of the document received from Alice. If the two values match, Bob can be sure that the document came from Alice and the document was not tampered with in transit. The slightest change in the document will cause the values to not match and will cause the authentication to fail.

**NOTE**

When the message digest generated is encrypted using a key, it's called a keyed message digest. Another definition for a keyed message digest is *message authentication code (HMAC)*.

## IPSec Security Protocols

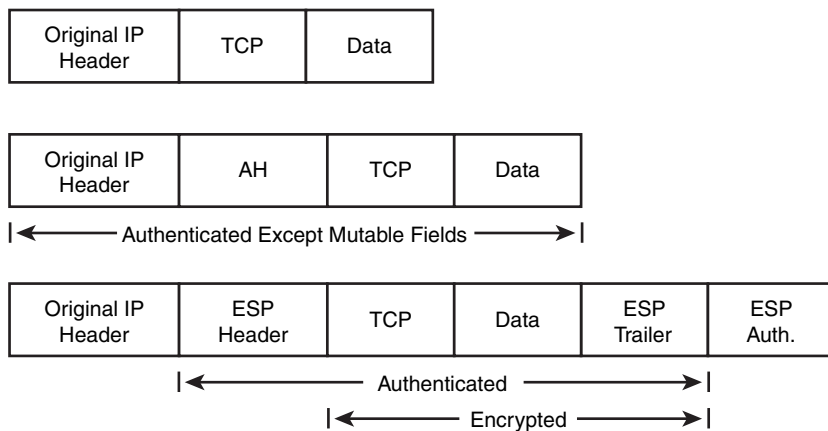
The objective of IPSec is to provide security services for IP packets at the network layer. These services include access control, data integrity, authentication, protection against replay, and data confidentiality.

Encapsulating security payload (ESP) and authentication header (AH) are the two IPSec security protocols used to provide this security for an IP datagram. Before looking into the IPSec security protocols, you must understand the two IPSec modes, transport and tunnel mode, and what services each provides.

## IPSec Transport Mode

In transport mode, an IPSec header (AH or ESP) is inserted between the IP header and the upper layer protocol header. Figure 2-3 shows an IP packet protected by IPSec in transport mode.

**Figure 2-3** *IP Packet in IPSec Transport Mode*



In this mode, the IP header is the same as that of the original IP packet except for the IP protocol field, which is changed to ESP (50) or AH (51), and the IP header checksum, which is recalculated. IPSec assumes the IP endpoints are reachable. In this mode, the destination IP address in the IP header is not changed by the source IPSec endpoint; therefore, this mode can only be used to protect packets in scenarios in which the IP endpoints and the IPSec endpoints are the same.

From an IPSec VPN point of view, this mode is most useful when traffic between two hosts must be protected, rather than when traffic moves from site-to-site, and each site has many hosts. The biggest challenge when using IPSec transport mode in the site-to-site model is the complexity involved in managing IPSec protection from any given host to all the possible peer hosts. Additionally, the two hosts' IP addresses must be routable across the entire IP routing path. Due to the complexities of building an IPSec transport mode VPN from host to host, the typical VPN will use a VPN gateway to protect all the hosts from one site to all the hosts at a peer site. A typical IPSec VPN deployment occurs between sites where each site has multiple hosts behind a VPN gateway and the IPSec tunnel endpoints serve as the VPN gateway routers. With the VPN gateway protecting a set of host IP addresses, the IPSec transport mode has

limited utility. IPSec transport mode can still be used for VPN connectivity if Generic Route Encapsulation (GRE) IP tunnels are used between the VPN gateways. The GRE tunnel endpoints serve as “host” endpoints. IPSec protects the GRE tunnel traffic in transport mode. Chapter 3, “Enhanced IPSec Features,” explores more about GRE and IPSec.

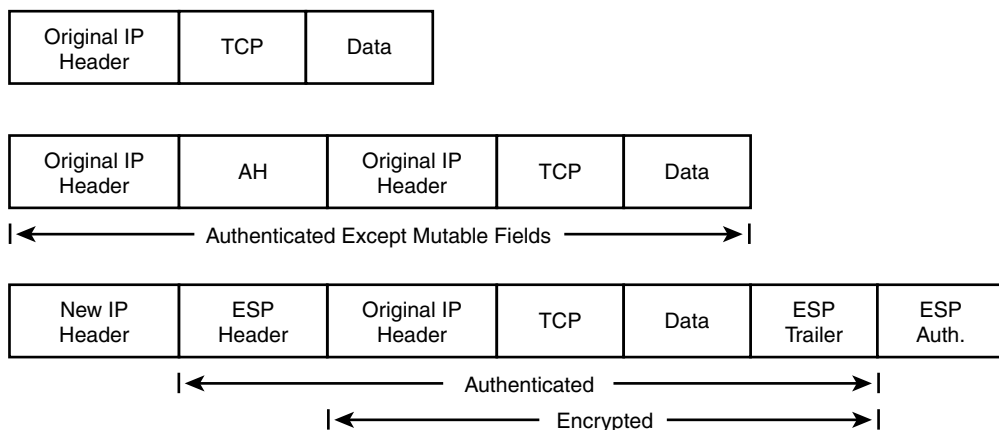
**NOTE** Another limitation of transport mode is that it cannot be used with NAT translation of packets between IPSec peers. Also, for most hardware encryption engines, it is less efficient to encrypt transport mode than tunnel mode, because transport mode requires displacement of the IP header to make room for the ESP or AH header.

## IPSec Tunnel Mode

IPSec VPN service using transport mode and GRE encapsulation between the VPN gateways at each site is a very popular option for site-to-site VPNs. But what about an IP node that has no GRE support, yet requires the establishment of IPSec VPN connectivity with another site? The most common example of this is a telecommuter. IPSec tunnel mode helps address this situation.

In tunnel mode, the original IP packet is encapsulated in another IP datagram, and an IPSec header (AH or ESP) is inserted between the outer and inner headers. Because of this encapsulation with an “outer” IP packet, tunnel mode can be used to provide security services between sites on behalf of IP nodes behind the gateway router at each site. Also, this mode can be used for the telecommuter scenario of connecting from an end host to an IPSec gateway at a site. Figure 2-4 shows an IP packet protected by IPSec in tunnel mode.

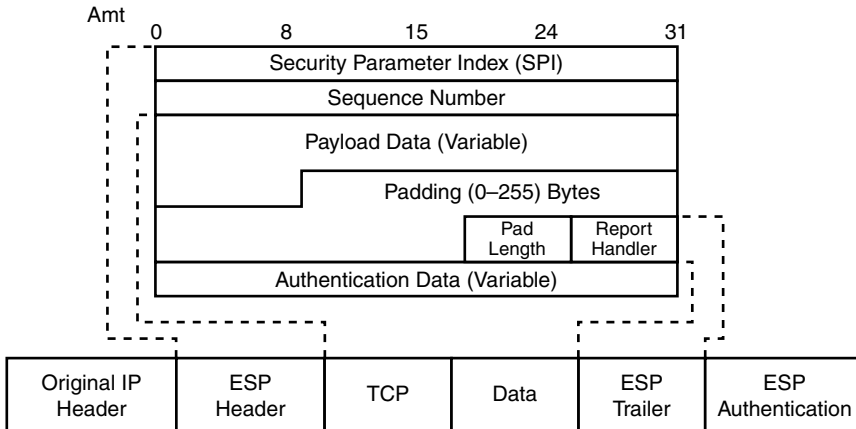
**Figure 2-4** *IP Packet in IPSec Tunnel Mode*



## Encapsulating Security Header (ESP)


ESP provides confidentiality, data integrity, and optional data origin authentication and anti-replay services. It provides these services by encrypting the original payload and encapsulating the packet between a header and a trailer, as shown in Figure 2-5.


**Figure 2-5** *IP Packet Protected by ESP*



ESP is identified by a value of 50 in the IP header. The ESP header is inserted after the IP header and before the upper layer protocol header. The IP header itself could be a new IP header in tunnel mode or the original IP packet's header in transport mode. Figures 2-6 and 2-7 show the position of the ESP header in transport and tunnel mode, respectively.

**Figure 2-6** *IP Packet Protected by ESP in Transport Mode*



**Figure 2-7** IP Packet Protected by ESP in Tunnel Mode

The *security parameter index (SPI)* in the ESP header is a 32-bit value that, combined with the destination address and protocol in the preceding IP header, identifies the security association (SA) to be used to process the packet. The SPI is an arbitrary number chosen by the destination peer during Internet Key Exchange (IKE) negotiation between the peers. It functions like an index number that can be used to look up the SA in the security association database (SADB).

The sequence number is a unique monotonically increasing number inserted into the header by the sender. Sequence numbers, along with the sliding receive window, provide anti-replay services. The anti-replay protection scheme is common to both ESP and AH.

The data being protected (or, more specifically, being encrypted by ESP) is in the payload data field. The algorithm used to encrypt the payload may require an initialization vector (IV), which is also carried in the data payload. Note that the IV is authenticated but not encrypted. If the encryption algorithm used is DES, then the first eight bytes of the protected data field is the IV; 3DES and AES also have an 8-byte IV.

*Padding* in the ESP header is the addition of bits to the ESP header; the number of bits to be padded depends on the encryption algorithm that is used. The Pad Length field indicates the number of pad bytes added so that the original data can be restored on decryption.

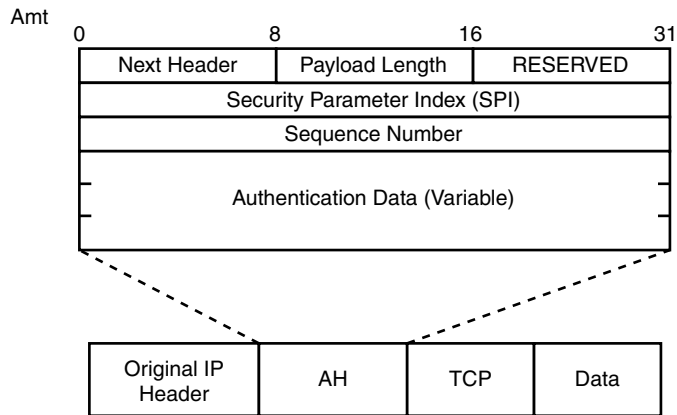
The next header payload identifies the type of data in the payload. For example, if ESP is used in tunnel mode, this value will be 4.

Authentication digest in the ESP header is used to verify data integrity. Because authentication is always applied after encryption, a check for validity of the data is done upon receipt of the packet and before decryption.

## Authentication Header (AH)


AH provides connectionless integrity, data authentication, and optional replay protection but, unlike ESP, it does not provide confidentiality. Consequently, it has a much simpler header than ESP. Figure 2-8 shows an AH-protected IP packet.

**Figure 2-8** *IP Packet Protected by AH*

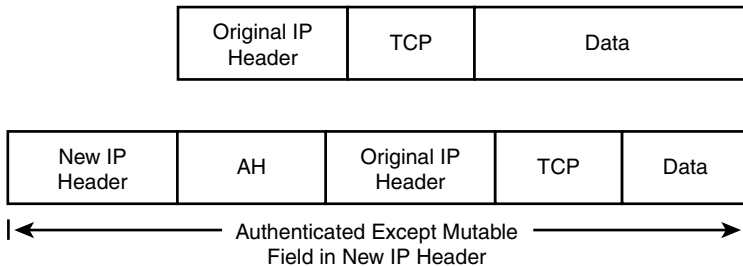


AH is an IP protocol, identified by a value of 51 in the IP header. The Next header field indicates what follows the AH header. In transport mode, it will be the value of the upper layer protocol being protected (for example, UDP or TCP). In tunnel mode, this value is 4. The positions of AH in transport and tunnel mode are shown in Figure 2-9 and Figure 2-10, respectively.

**Figure 2-9** *IP Packet Protected by AH in Transport Mode*



**Figure 2-10** *IP Packet Protected by AH in Tunnel Mode*



AH in transport mode is useful if the communication endpoints are also the IPSec endpoints. In tunnel mode, AH encapsulates the IP packet and an additional IP header is added before the AH header. Although the tunnel mode of AH could be used to provide IPSec VPN end-to-end security, there is no data confidentiality in AH and hence this mode is not too useful.

The payload length field in the AH header in Figure 2-9 indicates the length of the header. The Reserved field is not used, and is, therefore, set to zeroes. The SPI and sequence numbers have the same significance as in ESP. The authentication digest has one key difference from ESP: With AH, authentication is provided to the IP header in addition to the payload. As AH creates the authentication data on the entire packet, including the IP header, some of the IP fields will change in transit; therefore, all those fields in the IP header that may change in transit are zeroed out before the authentication digest is hashed. The fields that zero out include type of service (ToS) bits, flags, fragment offset, time-to-live (TTL), and header checksum. These fields are zeroed out because authenticating a changed value in transit (for example, TTL) will cause the authentication hash to have a mismatch from the sender and the packet will be dropped.

## Key Management and Security Associations

You learned that there are two types of encryption algorithms—symmetric and asymmetric. You also know that IPSec VPNs are typically deployed across a public infrastructure because IPSec offers encryption services to keep the data confidential from non-intended recipients of the data. DES and 3DES are two of the most popular encryption algorithms used for IPSec VPNs; both are symmetric algorithms and, therefore, have to deal with the challenge of secure key distribution. Problems arise when the key distribution must be done over a public infrastructure such as the Internet.

Collectively, the generation, distribution, and storage of keys is called key management. All crypto systems must deal with key management issues. The default IPSec method for secure key negotiation is the Internet Key Exchange (IKE) protocol. IKE is designed to provide mutual authentication of systems, as well as to establish a shared secret key to create IPSec security associations. Before delving into how IKE works, it may be helpful to review the Diffie-Hellman key management protocol that is used by IKE to exchange a secret key over an insecure medium (such as the Internet).

## The Diffie-Hellman Key Exchange

Whitfield Diffie and Martin Hellman first published their algorithm in 1976. This algorithm is based on the difficulty of solving the discrete logarithm problem. In short, the situation is as follows (using the classic cryptographic characters of Alice, Bob, and Eve):

- Alice wishes to communicate with Bob securely.
- In order to achieve this secure communication, Alice needs to establish a session key with Bob, but they have to somehow agree on a shared key over a public medium that is insecure.