



Virtual Private Networking with Windows Server 2003: Overview

Microsoft Corporation

Published: March 2003

Abstract

This white paper provides an overview of virtual private networking and the virtual private network (VPN) technologies supported by Windows Server 2003 and Windows XP. Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec) are described as the two industry standard methods for VPN connections. This paper also describes the set of features in Windows Server 2003 and Windows XP that provides advanced security capabilities and simplified administration of VPN connections for enterprise networks.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2003 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Windows, Windows NT, and Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Contents

Introduction	1
Common Uses of VPNs	2
Remote Access Over the Internet	2
Connecting Networks Over the Internet	2
Connecting Computers Over an Intranet	3
Basic VPN Requirements	4
Tunneling Basics	5
Tunneling Protocols	6
How Tunneling Works	6
Tunneling Protocols and the Basic VPN Requirements.....	6
Point-to-Point Protocol (PPP)	7
Phase 1: PPP Link Establishment.....	7
Phase 2: User Authentication.....	7
Phase 3: PPP Callback Control.....	9
Phase 4: Invoking Network Layer Protocol(s)	9
Data-Transfer Phase	9
Point-to-Point Tunneling Protocol (PPTP)	9
Layer Two Tunneling Protocol (L2TP)	10
PPTP Compared to L2TP/IPSec.....	11
Advantages of L2TP/IPSec Over PPTP	11
Advantages of PPTP Over L2TP/IPSec.....	12
Tunnel Types	12
Voluntary Tunneling	12
Compulsory Tunneling	13
Advanced VPN Security Features	14
EAP-TLS and Certificate-based Authentication.....	14
Digital Certificates	14
Extensible Authentication Protocol (EAP)	15
EAP-Transport Level Security (EAP-TLS)	15
Network Access Quarantine Control.....	15

Remote Access Account Lockout	16
Remote Access Policy Profile Packet Filtering	16
VPN Administration	18
Authorizing VPN Connections.....	18
Scalability	18
RADIUS.....	19
Connection Manager and Managed VPN Connections	19
Connection Manager Client Dialer	19
Connection Manager Administration Kit.....	20
Connection Point Services	20
Accounting, Auditing, and Alarming	22
Summary	23
Related Links	24

Introduction

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. A VPN enables you to send data between two computers across a shared or public internetwork in a manner that emulates the properties of a point-to-point private link. The act of configuring and creating a virtual private network is known as virtual private networking.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information allowing it to traverse the shared or public transit internetwork to reach its endpoint. To emulate a private link, the data being sent is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The portion of the connection in which the private data is encapsulated is known as the tunnel. The portion of the connection in which the private data is encrypted is known as the virtual private network (VPN) connection.




Figure 1: Virtual private network connection

VPN connections allow users working at home or on the road to connect in a secure fashion to a remote organization server using the routing infrastructure provided by a public internetwork (such as the Internet). From the user's perspective, the VPN connection is a point-to-point connection between the user's computer and an organization server. The nature of the intermediate internetwork is irrelevant to the user because it appears as if the data is being sent over a dedicated private link.

VPN technology also allows a corporation to connect to branch offices or to other companies over a public internetwork (such as the Internet), while maintaining secure communications. The VPN connection across the Internet logically operates as a wide area network (WAN) link between the sites.

In both of these cases, the secure connection across the internetwork appears to the user as a private network communication—despite the fact that this communication occurs over a public internetwork—hence, the name virtual private network.

VPN technology is designed to address issues surrounding the current business trend toward increased telecommuting and widely distributed global operations, where workers must be able to connect to central resources and must be able to communicate with one another.

To provide employees with the ability to connect to organization computing resources, regardless of their location, a corporation must deploy a scalable remote access solution. Typically, corporations choose either a department solution, where an internal information systems department is charged with buying, installing, and maintaining organization modem pools and a private network infrastructure; or

they choose a value-added network (VAN) solution, where they pay an outsourced company to buy, install, and maintain modem pools and a telecommunication infrastructure.

Neither of these solutions provides the necessary scalability, in terms of cost, flexible administration, and demand for connections. Therefore, it makes sense to replace the modem pools and private network infrastructure with a less expensive solution based on Internet technology so that the business can focus on its core competencies. With an Internet solution, a few Internet connections through Internet service providers (ISPs) and VPN server computers can serve the remote networking needs of hundreds or thousands of remote clients and branch offices.

Common Uses of VPNs

The next few sections describe the more common VPN configurations in more detail.

Remote Access Over the Internet

VPNs provide remote access to organization resources over the public Internet, while maintaining privacy of information. Figure 2 shows a VPN connection used to connect a remote access client to an organization intranet. This is known as a remote access VPN connection.




Figure 2: Using a VPN connection to connect a remote access client to an organization intranet

Rather than making a long distance (or 1-800) call to an organization or outsourced network access server (NAS), the user dials a local ISP. Using the connection to the local ISP, the VPN client creates a VPN connection between the remote access computer and the organization VPN server across the Internet.

Connecting Networks Over the Internet

There are two methods for using VPNs to connect local area networks at remote sites:

- Using dedicated lines to connect a branch office to an organization LAN.

For example, rather than using an expensive long-distance dedicated circuit between the branch office and the corporate hub, both the branch office and the corporate hub routers can use a local dedicated circuit and local ISP to connect to the Internet. The VPN software uses the local ISP connections and the Internet to create a virtual private network between the branch office router and corporate hub router.

- Using a dial-up line to connect a branch office to the Internet.

Rather than having a router at the branch office make a long distance (or 1-800) call to a corporate or outsourced NAS, the router at the branch office can call a local ISP. The VPN client uses the connection to the local ISP to create a VPN connection between the branch office router and the corporate hub router across the Internet. This is known as a site-to-site VPN connection.




Figure 3: Using a VPN connection to connect two remote sites

In both cases, the facilities that connect the branch office and corporate offices to the Internet are local. The corporate hub router that acts as a VPN server must be connected to a local ISP with a dedicated line. This VPN server must be listening 24 hours a day for incoming VPN traffic.

Connecting Computers Over an Intranet

In some organization internetworks, the departmental data is so sensitive that the department's LAN is physically disconnected from the rest of the organization internetwork. Although this protects the department's confidential information, it creates information accessibility problems for those users not physically connected to the separate LAN.




Figure 4: Using a VPN connection to connect to a secured or hidden network

VPNs allow the department's LAN to be physically connected to the organization internetwork but separated by a VPN server. The VPN server is not acting as a router between the organization internetwork and the department LAN. A router would connect the two networks, allowing everyone access to the sensitive LAN. By using a VPN server, the network administrator can ensure that only those users on the organization internetwork who have appropriate credentials (based on a need-to-know policy within the company) can establish a VPN connection with the VPN server and gain access to the protected resources of the department. Additionally, all communication across the VPN can be

encrypted for data confidentiality. Those users who do not have the proper credentials cannot view the department LAN.

Basic VPN Requirements

Typically, when deploying a remote networking solution, an enterprise needs to facilitate controlled access to organization resources and information. The solution must allow roaming or remote clients to connect to LAN resources, and the solution must allow remote offices to connect to each other to share resources and information (router-to-router connections). In addition, the solution must ensure the privacy and integrity of data as it traverses the Internet. The same concerns apply in the case of sensitive data traversing an organization internetwork.

Therefore, a VPN solution should provide at least all of the following:

- **User Authentication.** The solution must verify the VPN client's identity and restrict VPN access to authorized users only. It must also provide audit and accounting records to show who connected and for how long.
- **Address Management.** The solution must assign a VPN client an address on the intranet and ensure that addresses used on the intranet are kept private.
- **Data Encryption.** Data carried on the public network must be rendered unreadable.
- **Key Management.** The solution must generate and refresh encryption keys for the encrypted data.

An Internet VPN solution based on the Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPSec) meets all of these basic requirements and takes advantage of the broad availability of the Internet. Other solutions, including IPSec tunnel mode, meet only some of these requirements, but remain useful for specific situations.

The remainder of this paper discusses VPN concepts, protocols, and components in greater detail.

Tunneling Basics

Tunneling is a method of using an internetwork infrastructure to transfer data for one network over another network. The data to be transferred (or payload) can be the frames (or packets) of another protocol. Instead of sending a frame as it is produced by the originating node, the tunneling protocol encapsulates the frame in an additional header. The additional header provides routing information so that the encapsulated payload can traverse the intermediate internetwork.

The encapsulated packets are then routed between tunnel endpoints over the internetwork. The logical path through which the encapsulated packets travel through the internetwork is called a tunnel. Once the encapsulated frames reach their destination on the internetwork, the frame is decapsulated and forwarded to its final destination. Tunneling includes this entire process (encapsulation, transmission, and decapsulation of packets).




Figure 5: Tunneling

The transit internetwork can be any internetwork—the Internet is a public internetwork and is the most widely known real world example. There are many examples of tunnels that are carried over organization internetworks. And while the Internet provides one of the most pervasive and cost-effective internetworks, references to the Internet in this paper can be replaced by any other public or private internetwork that acts as a transit internetwork.

Tunneling technologies have been in existence for some time, such as SNA tunneling over IP internetworks. When System Network Architecture (SNA) traffic is sent across an organization IP internetwork, the SNA frame is encapsulated in a UDP and IP header. New tunneling technologies have been introduced in recent years. These newer technologies, which are the primary focus of this paper, include:

- **Point-to-Point Tunneling Protocol (PPTP).** PPTP allows multiprotocol traffic to be encrypted and then encapsulated in an IP header to be sent across an organization IP internetwork or a public IP internetwork such as the Internet.
- **Layer Two Tunneling Protocol (L2TP).** L2TP allows multiprotocol traffic to be encrypted and then sent over any medium that supports point-to-point datagram delivery, such as IP, X.25, Frame Relay, or ATM.
- **IPSec tunnel mode.** IPSec tunnel mode allows IP packets to be encrypted and then encapsulated in an IP header to be sent across an organization IP internetwork or a public IP internetwork such as the Internet. IPSec tunnel mode is not a recommended technology for remote access VPN connections, because there are no standard methods for user authentication, IP address assignment, and name

server address assignment. Using IPSec tunnel mode for site-to-site VPN connections is possible using computers running Windows Server 2003. Because the IPSec tunnel is not represented as a logical interface over which packets can be forwarded and received, routes cannot be assigned to use the IPSec tunnel and routing protocols do not operate over IPSec tunnels. Therefore, the use of IPSec tunnel mode is only recommended as a VPN solution for site-to-site VPN connections in which one end of the tunnel is a third-party VPN server or security gateway that does not support L2TP/IPSec.

Tunneling Protocols

For a tunnel to be established, both the tunnel client and the tunnel server must be using the same tunneling protocol. Tunneling technology can be based on either a Layer 2 or a Layer 3 tunneling protocol. These layers correspond to the Open Systems Interconnection (OSI) Reference Model. Layer 2 protocols correspond to the data-link layer and use frames as their unit of exchange. PPTP and L2TP are Layer 2 tunneling protocols; both encapsulate the payload in a PPP frame to be sent across an internetwork. Layer 3 protocols correspond to the Network layer, and use packets. IPSec tunnel mode is an example of a Layer 3 tunneling protocol and encapsulates IP packets with an additional IP header before sending them across an IP internetwork.

How Tunneling Works

For PPTP and L2TP, a tunnel is similar to a session; both of the tunnel endpoints must agree to the tunnel and must negotiate configuration variables, such as address assignment or encryption or compression parameters. In most cases, data transferred across the tunnel is sent using a datagram-based protocol. A tunnel management protocol is used as the mechanism to create, maintain, and terminate the tunnel.

Once the tunnel is established, tunneled data can be sent. The tunnel client or server uses a tunnel data transfer protocol to prepare the data for transfer. For example, when the tunnel client sends a payload to the tunnel server, the tunnel client first appends a tunnel data transfer protocol header to the payload. The client then sends the resulting encapsulated payload across the internetwork, which routes it to the tunnel server. The tunnel server accepts the packets, removes the tunnel data transfer protocol header, and forwards the payload to the target network. Information sent between the tunnel server and the tunnel client behaves similarly.

Tunneling Protocols and the Basic VPN Requirements

Because they are based on the well-defined PPP protocol, PPTP and L2TP inherit a suite of useful features. These features address the basic VPN requirements, as outlined below.

- **User Authentication.** PPTP and L2TP inherit the user authentication schemes of PPP, including the EAP methods discussed later in this paper. Using the Extensible Authentication Protocol (EAP), PPTP and L2TP connections can support a wide variety of authentication methods, including one-time passwords, cryptographic calculators, and smart cards.
- **Dynamic Address Assignment.** PPTP and L2TP connections support dynamic assignment of client addresses based on the Network Control Protocol (NCP) negotiation mechanism. For example, IP uses the Internet Protocol Control Protocol (IPCP) to negotiate an IP address and the addresses of name resolution servers.

- **Data Compression.** PPTP and L2TP support PPP-based compression schemes. For example, the Microsoft implementations of both PPTP and L2TP use Microsoft Point-to-Point Compression (MPPC).
- **Data Encryption.** PPTP and L2TP support PPP-based data encryption mechanisms. The Microsoft implementation of PPTP supports the use of Microsoft Point-to-Point Encryption (MPPE), based on the RSA/RC4 algorithm. The Microsoft implementation of L2TP uses IPsec encryption to protect the data stream from the VPN client to the VPN server.
- **Key Management.** MPPE for PPTP connections relies on the initial key generated during user authentication, and then refreshes it periodically. IPsec for L2TP/IPsec connections explicitly negotiates a common key during the IKE exchange, and also refreshes it periodically.

Point-to-Point Protocol (PPP)

Because PPTP and L2TP depend heavily on the features originally specified for PPP, it is worth examining this protocol more closely. PPP was designed to send data across dial-up or dedicated point-to-point connections. For IP, PPP encapsulates IP packets within PPP frames, and then transmits the PPP-encapsulated packets across a point-to-point link. PPP was originally defined as the protocol to use between a dial-up client and a NAS.

There are four distinct phases of negotiation in a PPP connection. Each of these four phases must complete successfully before the PPP connection is ready to transfer user data.

Phase 1: PPP Link Establishment

PPP uses the Link Control Protocol (LCP) to establish, maintain, and terminate the logical point-to-point connection. During Phase 1, basic communication options are selected. For example, authentication protocols are selected, but they are not actually implemented until the connection authentication phase (Phase 2). Similarly, during Phase 1, a decision is made as to whether the two peers will negotiate the use of compression and/or encryption. The actual choice of compression and encryption algorithms and other details occurs during Phase 4.

Phase 2: User Authentication

In the second phase, the client computer sends the user's credentials to the remote access server. A secure authentication scheme provides protection against replay attacks and remote client impersonation. A replay attack occurs when a third party monitors a successful connection and uses captured packets to play back the remote client's response so that it can gain an authenticated connection. Remote client impersonation occurs when a third party takes over an authenticated connection. The intruder waits until the connection has been authenticated and then traps the communication parameters, disconnects the authenticated user, and takes control of the authenticated connection.

Windows Server 2003 and Windows XP supports the following PPP authentication protocols:

- Password Authentication Protocol (PAP)

PAP is a simple, clear-text authentication scheme. The NAS requests the user name and password, and PAP returns them in clear text (unencrypted). Obviously, this authentication scheme is not secure because a third party could capture the user's name and password and use it to get

subsequent access to the NAS and all of the resources provided by the NAS. PAP provides no protection against replay attacks or remote client impersonation once the user's password is compromised.

- Challenge-Handshake Authentication Protocol (CHAP)

CHAP is an encrypted authentication mechanism that avoids transmission of the actual password on the connection. The NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must use the MD5 one-way hashing algorithm to return the user name and a hash of the challenge, session ID, and the client's password. The user name is sent as plain text.

CHAP is an improvement over PAP because the clear-text password is not sent over the link. Instead, the password is used to create a hash from the original challenge. The server knows the client's clear-text password and can, therefore, replicate the operation and compare the result to the password sent in the client's response. CHAP protects against replay attacks by using an arbitrary challenge string for each authentication attempt. CHAP protects against remote client impersonation by unpredictably sending repeated challenges to the remote client throughout the duration of the connection.

- Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP)

MS-CHAP is an encrypted authentication mechanism very similar to CHAP. As in CHAP, the NAS sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must return the user name and an encrypted form of the challenge string, the session ID, and the MD4-hashed password. This design, which uses the MD4 hash of the password, provides an additional level of security because it allows the server to store hashed passwords instead of clear-text passwords. MS-CHAP also provides additional error codes, including a password expired code, and additional encrypted client-server messages that permit users to change their passwords during the authentication process. In MS-CHAP, both the access client and the NAS independently generate an initial encryption key for subsequent data encryption by MPPE. Therefore, MS-CHAP authentication is required to enable MPPE-based data encryption.

- MS-CHAP version 2 (MS-CHAP v2)

MS-CHAP v2 is an updated encrypted authentication mechanism that provides stronger security for the exchange of user name and password credentials and determination of encryption keys. With MS-CHAP v2, the NAS sends a challenge to the access client that consists of a session identifier and an arbitrary challenge string. The remote access client sends a response that contains the user name, an arbitrary peer challenge string, and an encrypted form of the received challenge string, the peer challenge string, the session identifier, and the user's password. The NAS checks the response from the client and sends back a response containing an indication of the success or failure of the connection attempt and an authenticated response based on the sent challenge string, the peer challenge string, the encrypted response of the client, and the user's password. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

Using this process, MS-CHAP v2 provides mutual authentication—the NAS verifies that the access client has knowledge of the user's password and the access client verifies that the NAS has

knowledge of the user's password. MS-CHAP v2 also determines two encryption keys, one for data sent and one for data received.

- Extensible Authentication Protocol (EAP).

EAP is a new PPP authentication protocol that allows for an arbitrary authentication method. EAP is described in the "Extensible Authentication Protocol (EAP)" section of this paper. EAP differs from the other authentication protocols in that EAP during the authentication phase does not actually perform authentication. Phase 2 for EAP only negotiates the use of a common EAP authentication method (known as an EAP type). The actual authentication for the negotiated EAP type is performed after Phase 2.

During phase 2 of PPP link configuration, the NAS collects the authentication data and then validates the data against its own user database or a central authentication database server, such as one maintained by a Windows domain controller, or the authentication data is sent to a Remote Authentication Dial-in User Service (RADIUS) server.

Phase 3: PPP Callback Control

The Microsoft implementation of PPP includes an optional callback control phase. This phase uses the Callback Control Protocol (CBCP) immediately after the authentication phase. If configured for callback, both the remote client and NAS disconnect after authentication. The NAS then calls the remote client back at a specified phone number. This provides an additional level of security to dial-up connections. The NAS allows connections from remote clients physically residing at specific phone numbers only. Callback is only used for dial-up connections, not for VPN connections.

Phase 4: Invoking Network Layer Protocol(s)

Once the previous phases have been completed, PPP invokes the various network control protocols (NCPs) that were selected during the link establishment phase (Phase 1) to configure protocols used by the remote client. For example, during this phase, IPCP is used to assign a dynamic address to the PPP client. In the Microsoft implementation of PPP, the Compression Control Protocol (CCP) is used to negotiate both data compression (using MPPE) and data encryption (using MPPE).

Data-Transfer Phase

Once the four phases of PPP negotiation have been completed, PPP begins to forward data to and from the two peers. Each transmitted data packet is wrapped in a PPP header that is removed by the receiving system. If data compression was selected in phase 1 and negotiated in phase 4, data is compressed before transmission. If data encryption is selected and negotiated, data is encrypted before transmission. If both encryption and compression are negotiated, the data is compressed first, and then encrypted.

Point-to-Point Tunneling Protocol (PPTP)

PPTP encapsulates PPP frames in IP datagrams for transmission over an IP internetwork, such as the Internet. PPTP can be used for remote access and router-to-router VPN connections. PPTP is documented in RFC 2637.

The Point-to-Point Tunneling Protocol (PPTP) uses a TCP connection for tunnel management and a modified version of Generic Routing Encapsulation (GRE) to encapsulate PPP frames for tunneled data. The payloads of the encapsulated PPP frames can be encrypted and/or compressed. Figure 6 shows the structure of a PPTP packet containing an IP datagram.




Figure 6. Structure of a PPTP packet containing an IP datagram

Layer Two Tunneling Protocol (L2TP)

L2TP is a combination of PPTP and Layer 2 Forwarding (L2F), a technology proposed by Cisco Systems, Inc. L2TP represents the best features of PPTP and L2F. L2TP encapsulates PPP frames to be sent over IP, X.25, Frame Relay, or Asynchronous Transfer Mode (ATM) networks. When configured to use IP as its datagram transport, L2TP can be used as a tunneling protocol over the Internet. L2TP is documented in RFC 2661.

L2TP over IP internetworks uses UDP and a series of L2TP messages for tunnel management. L2TP also uses UDP to send L2TP-encapsulated PPP frames as the tunneled data. The payloads of encapsulated PPP frames can be encrypted and/or compressed, although the Microsoft implementation of L2TP does not use MPPE to encrypt the PPP payload. Figure 7 shows the structure of an L2TP packet containing an IP datagram.

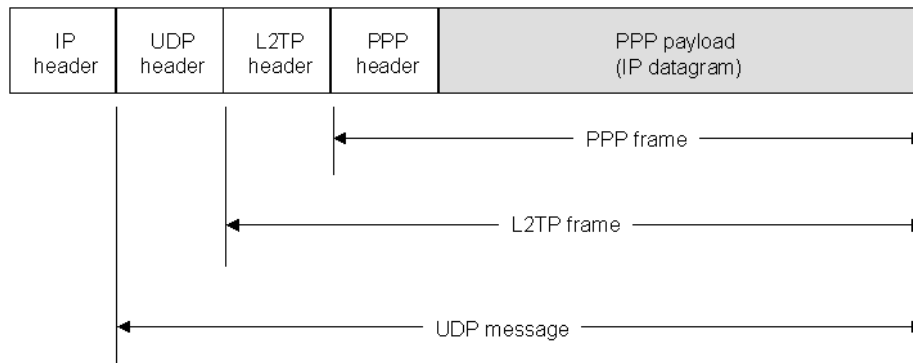


Figure 7. Structure of an L2TP packet containing an IP datagram

In the Microsoft implementation of L2TP, IPSec Encapsulating Security Payload (ESP) is used to encrypt L2TP traffic. The combination of L2TP (the tunneling protocol) and IPSec (the method of encryption) is known as L2TP/IPSec. L2TP/IPSec is described in RFC 3193.

The result after applying ESP to an IP packet containing an L2TP message is shown in Figure 8.

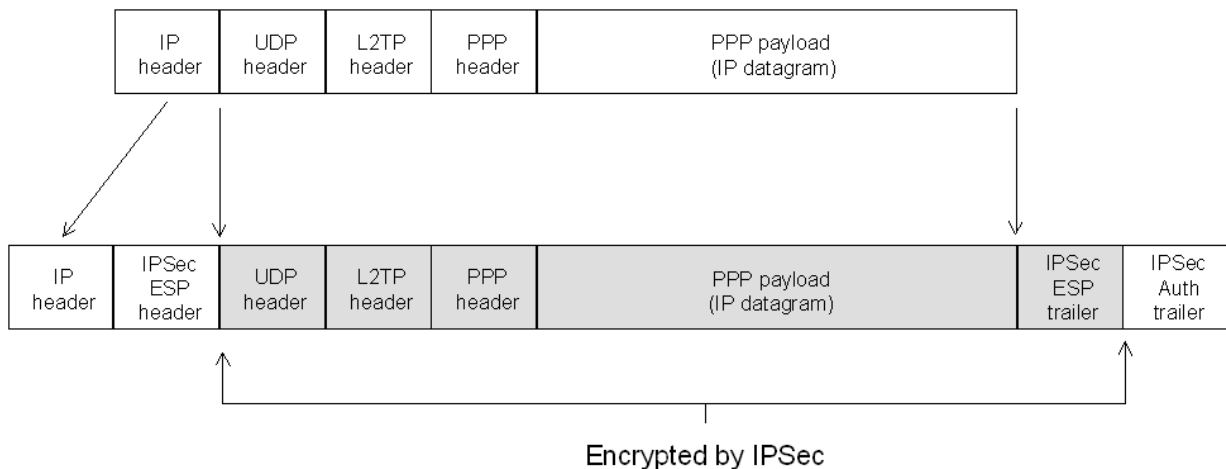


Figure 8. Encryption of L2TP traffic with IPsec ESP

PPTP Compared to L2TP/IPsec

Both PPTP and L2TP/IPsec use PPP to provide an initial envelope for the data, and then append additional headers for transport through the internetwork. However, there are the following differences:

- With PPTP, data encryption begins after the PPP connection process (and, therefore, PPP authentication) is completed. With L2TP/IPsec, data encryption begins before the PPP connection process by negotiating an IPsec security association.
- PPTP connections use MPPE, a stream cipher that is based on the Rivest-Shamir-Aldeman (RSA) RC-4 encryption algorithm and uses 40, 56, or 128-bit encryption keys. Stream ciphers encrypt data as a bit stream. L2TP/IPsec connections use the Data Encryption Standard (DES), which is a block cipher that uses either a 56-bit key for DES or three 56-bit keys for 3-DES. Block ciphers encrypt data in discrete blocks (64-bit blocks, in the case of DES).
- PPTP connections require only user-level authentication through a PPP-based authentication protocol. L2TP/IPsec connections require the same user-level authentication and, in addition, computer-level authentication using computer certificates.

Advantages of L2TP/IPsec Over PPTP

The following are the advantages of using L2TP/IPsec over PPTP in Windows Server 2003:

- IPsec ESP provides per-packet data origin authentication (proof that the data was sent by the authorized user), data integrity (proof that the data was not modified in transit), replay protection (prevention from resending a stream of captured packets), and data confidentiality (also known as encryption, providing prevention from interpreting captured packets without the encryption key). By contrast, PPTP provides only per-packet data confidentiality.
- L2TP/IPsec connections provide stronger authentication by requiring both computer-level authentication through certificates and user-level authentication through a PPP authentication protocol.

- PPP packets exchanged during user-level authentication are never sent in an unencrypted form because the PPP connection process for L2TP/IPSec occurs after the IPSec security association is established. If intercepted, the PPP authentication exchange for some types of PPP authentication protocols can be used to perform offline dictionary attacks and determine user passwords. By encrypting the PPP authentication exchange, offline dictionary attacks are much more difficult, as the encrypted packets must first be successfully decrypted.

Advantages of PPTP Over L2TP/IPSec

The following are advantages of PPTP over L2TP/IPSec in Windows Server 2003:

- PPTP does not require a certificate infrastructure. L2TP/IPSec requires a certificate infrastructure for issuing computer certificates to the VPN server computer and all VPN client computers.
- PPTP clients can be placed behind a network address translator (NAT) if the NAT has an editor for PPTP traffic. L2TP/IPSec-based VPN clients or servers cannot be placed behind a NAT unless both the VPN client and server support IPSec NAT traversal (NAT-T). IPSec NAT-T is supported by Windows Server 2003, Microsoft L2TP/IPSec VPN Client (for details, see “Microsoft L2TP/IPSec VPN Client” at <http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/l2tpclient.asp>), and L2TP/IPSec NAT-T Update for Windows XP and Windows 2000 (for details, see “L2TP/IPSec NAT-T Update for Windows XP and Windows 2000” at <http://support.microsoft.com/default.aspx?scid=kb;en-us;818043>).

Tunnel Types

Tunnels can be created in various ways.

- Voluntary tunnels

A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the user’s computer is a tunnel endpoint and acts as the tunnel client.

- Compulsory tunnels

A VPN-capable dial-up access server configures and creates a compulsory tunnel. With a compulsory tunnel, the user’s computer is not a tunnel endpoint. Another device, the dial-up access server, between the user’s computer and the tunnel server is the tunnel endpoint and acts as the tunnel client.

To date, voluntary tunnels are proving to be the more popular type of tunnel. The following sections describe each of these tunnel types in greater detail.

Voluntary Tunneling

Voluntary tunneling occurs when a workstation or routing server uses tunneling client software to create a virtual connection to the target tunnel server. To accomplish this, the appropriate tunneling protocol must be installed on the client computer. For the protocols discussed in this paper, voluntary tunnels require an IP connection (either LAN or dial-up).

In a dial-up situation, the client must establish a dial-up connection to the internetwork before the client can set up a tunnel. This is the most common case. The best example of this is the dial-up Internet

user, who must dial an ISP and obtain an Internet connection before a tunnel over the Internet can be created.

For a LAN-attached computer, the client already has a connection to the internetwork that can provide routing of encapsulated payloads to the chosen LAN tunnel server. This would be the case for a client on an organization LAN that initiates a tunnel to reach a private or hidden subnet on that LAN (such as the Human Resources network discussed previously).

It is a common misconception that VPN connections require a dial-up connection. They require only IP connectivity between the VPN client and VPN server. Some clients (such as home computers) use dial-up connections to the Internet to establish IP transport. This is a preliminary step in preparation for creating a tunnel and is not part of the tunnel protocol itself.

Compulsory Tunneling

A number of vendors that sell dial-up access servers have implemented the ability to create a tunnel on behalf of a dial-up client. The computer or network device providing the tunnel for the client computer is variously known as a Front End Processor (FEP) for PPTP or an L2TP Access Concentrator (LAC) for L2TP. For the purposes of this white paper, the term FEP is used to describe this functionality, regardless of the tunneling protocol. To carry out its function, the FEP must have the appropriate tunneling protocol installed and must be capable of establishing the tunnel when the client computer connects.

In the Internet example, the client computer places a dial-up call to a tunneling-enabled NAS at the ISP. For example, a corporation may have contracted with an ISP to deploy a nationwide set of FEPs. These FEPs can establish tunnels across the Internet to a tunnel server connected to the organization's private network, thus consolidating calls from geographically diverse locations into a single Internet connection at the organization network.

This configuration is known as compulsory tunneling because the client is compelled to use the tunnel created by the FEP. Once the initial connection is made, all network traffic to and from the client is automatically sent through the tunnel. With compulsory tunneling, the client computer makes a single PPP connection. When a client dials into the NAS, a tunnel is created and all traffic is automatically routed through the tunnel. An FEP can be configured to tunnel all dial-up clients to a specific tunnel server. The FEP could also tunnel individual clients, based on the user name or destination.

Unlike the separate tunnels created for each voluntary client, multiple dial-up clients can share a tunnel between the FEP and the tunnel server. When a second client dials into the access server (FEP) to reach a destination for which a tunnel already exists, there is no need to create a new instance of the tunnel between the FEP and tunnel server. Instead, the data traffic for the new client is carried over the existing tunnel. Since there can be multiple clients in a single tunnel, the tunnel is not terminated until the last user of the tunnel disconnects.

Advanced VPN Security Features

Because the Internet facilitates the creation of VPN connections from anywhere, networks need strong security features to prevent unwelcome access to private networks and to protect private data as it traverses the public network. User authentication and data encryption have already been discussed. This section provides overviews of advanced security features that can be used with Windows Server 2003 and Windows XP VPN connections.

EAP-TLS and Certificate-based Authentication

Symmetric, or private-key, encryption (also known as conventional encryption) is based on a secret key that is shared by both communicating parties. The sending party uses the secret key as part of the mathematical operation to encrypt (or encipher) plain text to cipher text. The receiving party uses the same secret key to decrypt (or decipher) the cipher text to plain text. Examples of symmetric encryption schemes are the RSA RC4 algorithm, which provides the basis for Microsoft Point-to-Point Encryption (MPPE), and Data Encryption Standard (DES), which is used for IPSec encryption.

Asymmetric, or public-key, encryption uses two different keys for each user: one is a private key known only to this one user; the other is a corresponding public key, which is accessible to anyone. The private and public keys are mathematically related by the encryption algorithm. One key is used for encryption and the other for decryption, depending on the nature of the communication service being implemented.

In addition, public key encryption technologies allow digital signatures to be placed on messages. A digital signature uses the sender's private key to encrypt some portion of the message. When the message is received, the receiver uses the sender's public key to decipher the digital signature to verify the sender's identity.

Digital Certificates

With symmetric encryption, both sender and receiver have a shared secret key. The distribution of the secret key must occur (with adequate protection) prior to any encrypted communication. However, with asymmetric encryption, the sender uses a private key to encrypt or digitally sign messages, while the receiver uses a public key to decipher these messages. The public key can be freely distributed to anyone who needs to receive the encrypted or digitally signed messages. The sender needs to carefully protect the private key only.

To secure the integrity of the public key, the public key is published with a certificate. A certificate (or public key certificate) is a data structure that is digitally signed by a certification authority (CA)—an authority that users of the certificate can trust. The certificate contains a series of values, such as the certificate name and usage, information identifying the owner of the public key, the public key itself, an expiration date, and the name of the certificate authority. The CA uses its private key to sign the certificate. If the receiver knows the public key of the certificate authority, the receiver can verify that the certificate is indeed from the trusted CA and, therefore, contains reliable information and a valid public key. Certificates can be distributed electronically (through Web access or email), on smart cards, or on floppy disks.

In summary, public key certificates provide a convenient, reliable method for verifying the identity of a sender. IPSec can optionally use this method for peer-level authentication. Remote access servers can use public key certificates for user authentication, as described in the section "EAP-Transport Level Security (EAP-TLS)."

Extensible Authentication Protocol (EAP)

As stated previously, most implementations of PPP provide very limited authentication methods. EAP is an IETF standard extension to PPP that allows for arbitrary authentication mechanisms for the validation of a PPP connection. EAP was designed to allow the dynamic addition of authentication plug-in modules at both the client and server ends of a connection. This allows vendors to supply a new authentication scheme at any time. EAP provides the highest flexibility in authentication uniqueness and variation.

EAP is documented in RFC 2284 and is supported in Windows Server 2003 and Windows XP.

EAP-Transport Level Security (EAP-TLS)

EAP-TLS is an IETF standard (RFC 2716) for a strong authentication method based on public-key certificates. With EAP-TLS, a client presents a user certificate to the dial-in server, and the server presents a server certificate to the client. The first provides strong user authentication to the server; the second provides assurance that the user has reached the server that he or she expected. Both systems rely on a chain of trusted authorities to verify the validity of the offered certificate.

The user's certificate could be stored on the VPN client computer or in an external smart card. In either case, the certificate cannot be accessed without some form of user identification (PIN number or name-and-password exchange) between the user and the client computer. This approach meets the something-you-know-plus-something-you-have criteria recommended by most security experts.

EAP-TLS is supported in Windows Server 2003 and Windows XP. Like MS-CHAP and MS-CHAP v2, EAP-TLS returns an encryption key to enable subsequent data encryption by MPPE.

Network Access Quarantine Control

Network Access Quarantine Control, a new feature in the Windows Server 2003 family, delays normal remote access to a private network until the configuration of the remote access computer has been examined and validated by an administrator-provided script. When a remote access computer initiates a connection to a remote access server, the user is authenticated and the remote access computer is assigned an IP address. However, the connection is placed in quarantine mode, with which network access is limited. The administrator-provided script is run on the remote access computer. When the script completes successfully, it runs a notifier component that notifies the remote access server that the remote access computer complies with current network policies. The remote access server removes quarantine mode, and the remote access computer is granted normal remote access.

Network Access Quarantine Control is a combination of the following:

- A remote access server running Windows Server 2003 and a quarantine notification listener service
- A RADIUS server running Windows Server 2003 and Internet Authentication Service (IAS), configured with a quarantine remote access policy that specifies quarantine settings

- A Connection Manager profile created with the Windows Server 2003 Connection Manager Administration Kit that contains a network policy compliance script and a notifier component
- A remote access client that is running Windows Server 2003, Windows XP, Windows 2000, Windows Millennium Edition, or Windows 98 Second Edition

For more information, see Windows Server 2003 Help and Support and the "Windows Server 2003 Network Access Quarantine Control" white paper on the [Windows VPN Web site](#).

Remote Access Account Lockout

The remote access account lockout feature is used to specify how many times a remote access authentication fails against a valid user account before the user is denied remote access. Remote access account lockout is especially important for remote access VPN connections over the Internet. Malicious users on the Internet can attempt to access an organization intranet by sending credentials (valid user name, guessed password) during the VPN connection authentication process. During a dictionary attack, the malicious user sends hundreds or thousands of credentials by using a list of passwords based on common words or phrases. With remote access account lockout enabled, a dictionary attack is thwarted after a specified number of failed attempts.

The remote access account lockout feature does not distinguish between malicious users who attempt to access your intranet and authentic users who attempt remote access but have forgotten their current passwords. Users who have forgotten their current password typically try the passwords that they remember and might have their accounts locked out.

If you enable the remote access account lockout feature, a malicious user can deliberately force an account to be locked out by attempting multiple authentications with the user account until the account is locked out, thereby preventing the authentic user from being able to log on.

The remote access account lockout feature is configured by changing settings in the registry on the computer that provides the authentication. If the remote access server is configured for Windows authentication, modify the registry on the remote access server computer. If the remote access server is configured for RADIUS authentication and Internet Authentication Service (IAS) is being used, modify the registry on the IAS server computer. For more information, see the topic titled "Remote access account lockout" in Windows Server 2003 Help.

Note: The remote access account lockout feature is not related to the **Account locked out** setting on the **Account** tab on the properties of a user account and the administration of account lockout policies using Group Policy.

Remote Access Policy Profile Packet Filtering

Remote access policies that define authorization and connection constraints can be used to specify a set of IP packet filters that are applied to remote access connections. When the connection is accepted, the packet filters define the types of IP traffic that are allowed from and to the VPN client.

This feature can be used for extranet connections. An extranet is a portion of your organization network that is accessible to users outside the organization, such as business partners and vendors. By using remote access policy profile packet filtering, you can create a remote access policy that specifies that

members of the Partners group can only access the Web servers at specific IP addresses or on a specific subnet.

This feature can also be used to prevent VPN remote access clients from sending packets that they did not originate. When the remote access client computer makes the VPN connection, by default it creates a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the remote access VPN client, treating the remote access client computer as a router, then that traffic is also be forwarded across the VPN connection. This is a security problem because the VPN server has not authenticated the computer that is forwarding traffic to the remote access VPN client. The computer forwarding traffic to the remote access VPN client computer has the same network access as the authenticated remote access VPN client computer.

To prevent the VPN server from receiving traffic across the VPN connection for computers other than authenticated remote access VPN client computers, configure remote access policy packet filters on the remote access policy that is used for your VPN connections. The default remote access policy for Windows Server 2003 named **Connections to Microsoft Routing and Remote Access server** already has the correct input packet filters for this configuration.

VPN Administration

In selecting a VPN technology, it is important to consider administrative issues. Large networks need to store per-user directory information in a centralized data store, or directory service, so that administrators and applications can add to, modify, or query this information. Each access or tunnel server could maintain its own internal database of per-user properties, such as names, passwords, and dial-in permission attributes. However, because it is administratively prohibitive to maintain multiple user accounts on multiple servers and keep them simultaneously current, most administrators set up an account database at the directory server or primary domain controller, or on a RADIUS server. By using the Active Directory® directory service as your account database, Windows Server 2003 VPNs become a single sign-on solution: the same set of credentials are used for both VPN connections to log on to the organization's domain.

Authorizing VPN Connections

To provide authorization for VPN connections and to provide a method of enforcing connection restraints, Windows Server 2003 VPN connections use a combination of the dial-in properties of user accounts in a local or domain account database and remote access policies.

Remote access policies are an ordered set of rules that define how connections are either accepted or rejected. For connections that are accepted, remote access policies can also define connection restrictions. For each rule, there are one or more conditions, a set of profile settings, and a remote access permission setting. Connection attempts are evaluated against the remote access policies in order, trying to determine whether the connection attempt matches all of the conditions of each policy. If the connection attempt does not match all of the conditions of any policy, the connection attempt is rejected.

If a connection matches all the conditions of a remote access policy and is granted remote access permission, the remote access policy profile specifies a set of connection restrictions. The dial-in properties of the user account also provide a set of restrictions. Where applicable, user account connection restrictions override the remote access policy profile connection restrictions. Remote access policy profile restrictions include connection settings (such as maximum connection time or an idle timeout), IP packet filtering, required authentication protocols, and required encryption strengths.

Scalability

Redundancy and load balancing is accomplished using either DNS or Network Load Balancing:

- Round-robin DNS is used to split requests among a number of VPN servers that share a common security perimeter. A security perimeter has one external DNS name—for example, microsoft.com—but several IP addresses, and loads are randomly distributed across all of the IP addresses.
- With Network Load Balancing, a cluster of VPN server computers can provide high availability and load balancing for both PPTP and L2TP/IPSec connections.

RADIUS

The Remote Authentication Dial-in User Service (RADIUS) protocol is a popular method for managing remote user authentication and authorization. RADIUS is a lightweight, UDP-based protocol. RADIUS servers can be located anywhere on the Internet and provide authentication (including PPP PAP, CHAP, MS-CHAP, MS-CHAP v2, and EAP) and authorization for access servers such as NASes and VPN servers.

In addition, RADIUS servers can provide a proxy service to forward authentication requests to distant RADIUS servers. For example, many ISPs have agreements to allow roaming subscribers to use local services from the nearest ISP for dial-up access to the Internet. These roaming alliances take advantage of the RADIUS proxy service. If an ISP recognizes a user name as being a subscriber to a remote network, the ISP uses a RADIUS proxy to forward the access request to the appropriate network.

Windows Server 2003 includes a RADIUS server and proxy with the Internet Authentication Service (IAS); an optional Windows networking component installed using Control Panel-Network.

Connection Manager and Managed VPN Connections

To deploy the configuration of a large number of VPN remote access clients for enterprise or outsourced dial scenarios, use Connection Manager (CM). CM is a set of components included with Windows Server 2003 that consists of the following:

- Connection Manager (CM) client dialer
- Connection Manager Administration Kit (CMAK)
- Connection Point Services (CPS)

Connection Manager Client Dialer

The Connection Manager (CM) client dialer is software that is installed on each VPN client. It includes advanced features that make it a superset of basic remote access networking. At the same time, CM presents a simplified dialing experience to the user. It limits the number of configuration options that a user can change, ensuring that the user can always connect successfully. For example, with the CM client dialer, a user can:

- Select from a list of phone numbers to use, based on physical location (for an outsourced VPN solution).
- Use customized graphics, icons, messages, and help.
- Automatically create a dial-up connection before the VPN connection is made.
- Run custom actions during various parts of the connection process, such as pre-connect and post-connect actions (executed before or after the dial-up or VPN connection is completed).

A customized CM client dialer package, also known as a profile, is a self-extracting executable file that is created by a network administrator with the Connection Manager Administration Kit (CMAK). The CM profile is distributed to VPN users via CD-ROM, e-mail, Web site, or file share. When the user runs the CM profile, it automatically configures the appropriate dial-up and VPN connections. The Connection Manager profile does not require a specific version of Windows. It will configure connections for

computers running Windows Server 2003, Windows XP, Windows 2000, Windows NT® 4.0, Windows Millennium Edition, and Windows 98.

Connection Manager Administration Kit

The Connection Manager Administration Kit (CMAK) is an optional management tool installed from:

- **Add/Remove Programs (in Control Panel) on a computer running Windows Server 2003.** You must specify Connection Manager Components in the Management and Monitoring Tools category of Windows components.
- **Windows Server 2003 Administration Tools on a computer running Windows XP Professional.** You must run the Adminpak.msi file from the \i386 folder on a Windows Server 2003 CD-ROM. After it is installed, you can run Connection Manager Administration Kit from Administrative Tools.

CMAK is a Wizard that guides you through a variety of options when configuring a CM profile and creates the profile to distribute to your VPN users.

Connection Point Services

Connection Point Services (CPS) allows you to create, distribute, and update custom phone books. Phone books contain one or more Point of Presence (POP) entries. Each POP has a telephone number used to access a dial-up network or the Internet. Phone books give users complete POP information, so when they travel, they can connect to different organization or Internet access points based on location, rather than having to use a toll-free or long distance number.

Without the ability to update phone books, users would not only have to contact their organization's technical support staff to obtain changes in POP information, they would also have to reconfigure their client dialer software.

CPS is a combination of:

- **Phone Book Administrator.** A tool used to both create and maintain phone book files and publish new or updated phone book files on the phone book server.
- **Phone Book Server.** A computer running Windows Server 2003 and Internet Information Services (IIS) (including the FTP Publishing Service) and an Internet Server Application Programming Interface (ISAPI) extension that processes phone book update requests from CM clients.

The Phone Book Administrator is a tool that is installed by running Pbainst.exe from the Valueadd\Msft\Mgmt\Pba folder on the Windows Server 2003 product CD-ROM. Once installed, you can run Phone Book Administrator from Administrative Tools. It is not required to run the Phone Book Administrator on the phone book server.

You can use the Phone Book Administrator to create phone book entries and regions and publish them in the *SystemRoot\Program Files\Phone Book Service\Data\PhoneBookFileName* folder of the phone book server.

After the phone book is configured and published, the CM profile is created with CMAK and configured with:

- Automatically downloaded phone book updates.

- The phone book file.
- The name of the phone book server.

Accounting, Auditing, and Alarming

To properly administer a VPN system, network administrators should be able to track who uses the system, how many connections are made, unusual activity, error conditions, and situations that may indicate equipment failure. This information can be used for billing, auditing, and alarm or error-notification purposes.

For example, an administrator may need to know who connected to the system and for how long in order to construct billing data. Unusual activity may indicate a misuse of the system or inadequate system resources. Real-time monitoring of equipment (for example, unusually high activity on one modem and inactivity on another) may generate alerts to notify the administrator of a modem failure. The tunnel server should provide all of this information, and the system should provide event logs, reports, and a data storage facility to handle the data appropriately.

The RADIUS protocol defines a suite of call-accounting requests that are independent from the authentication requests discussed above. These messages from the NAS to the RADIUS server request the latter to generate accounting records at the start of a call, the end of a call, and at predetermined intervals during a call. The Routing and Remote Access service, which provides the VPN server functionality in Windows Server 2003, can be configured to generate these RADIUS accounting requests separately from connection requests (which could go to the domain controller or to a RADIUS server). This allows an administrator to configure an accounting RADIUS server, whether RADIUS is used for authentication or not. An accounting server can then collect records for every VPN connection for later analysis. A number of third parties have already written billing and audit packages that read these RADIUS accounting records and produce various useful reports.

IAS in Windows Server 2003 is a RADIUS accounting server and supports recording the connection accounting information to a log file or sending it directly to a Structured Query Language (SQL) server database.

Summary

VPNs allow users or corporations to connect to remote servers, branch offices, or to other companies over a public internetwork, while maintaining secure communications. In all of these cases, the secure connection appears to the user as a private network communication—despite the fact that this communication occurs over a public internetwork. VPN technology is designed to address issues surrounding the current business trend toward increased telecommuting and widely distributed global operations, where workers must be able to connect to central resources and communicate with one another.

Virtual private networking with Windows Server 2003 and Windows XP supports the industry standard PPTP and L2TP/IPSec VPN protocols, advanced security features such as certificate-based authentication and Network Access Quarantine Control, and administration features such as centralized authentication and accounting with RADIUS and managed VPN client deployment with Connection Manager.

Related Links

See the following resources for further information:

- [Windows VPN Web site](http://www.microsoft.com/vpn) at <http://www.microsoft.com/vpn>
- [Microsoft L2TP/IPSec VPN Client](http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/l2tpclient.asp) at <http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/l2tpclient.asp>
- [Internet Authentication Service Web site](http://www.microsoft.com/windows2000/technologies/communications/ias/default.asp) at <http://www.microsoft.com/windows2000/technologies/communications/ias/default.asp>

For the latest information about Windows Server 2003, see the [Windows Server 2003 Web site](http://www.microsoft.com/windowsserver2003) at <http://www.microsoft.com/windowsserver2003>.