Operating System

Virtual Private Networking with Windows 2000: Deploying Remote Access VPNs

Microsoft Corporation
Published: July 2002

Abstract

A virtual private network (VPN) is the extension of a private network that encompasses logical links across shared or public networks such as the Internet. A remote access VPN connection allows computers connected to the Internet to securely access organization intranets. This paper describes the various components and design choices of a deployment of remote access VPN connections using the Windows® 2000 platform VPN servers and Windows-based VPN clients. This paper also includes detailed walkthroughs to deploy Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP)-based remote access VPNs, information on firewall configuration, how to create a VPN test lab, and details of troubleshooting tools and common problems. This paper assumes familiarity with TCP/IP, IP routing, Internet Protocol security (IPSec), and the capabilities of the Windows 2000 Routing and Remote Access service.

Contents

Introduction

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link (such as a long haul T-Carrier-based wide area network [WAN] link). Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a virtual private network (VPN) connection.

Figure 1 shows the logical equivalent of a VPN connection.

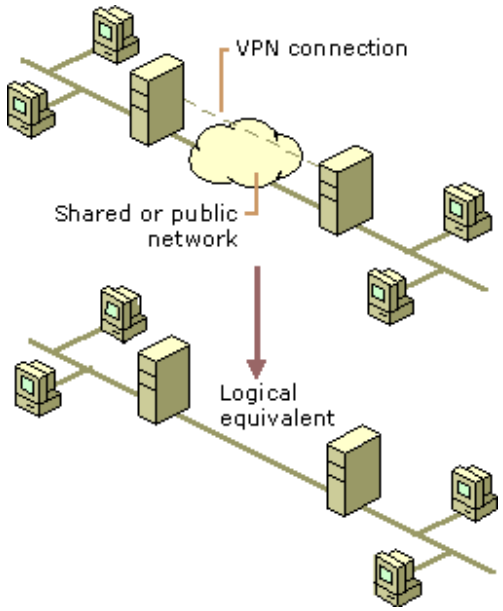

Figure 1 The logical equivalent of VPN connections

Users working at home or on the road can use VPN connections to establish a remote access connection to an organization server by using the infrastructure provided by a public network such as the Internet. From the user's perspective, the VPN connection is a point-to-point connection between the computer (the VPN client) and an organization server (the VPN server). The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations can also use VPN connections to establish routed connections with geographically separate offices or with other organizations over a public network such as the Internet while maintaining secure communications. A routed VPN connection across the Internet logically operates as a dedicated WAN link.

With both remote access and routed connections, an organization can use VPN connections to trade long-distance dial-up or leased lines for local dial-up or leased lines to an Internet service provider (ISP).

There are two types of remote access VPN technology in the Windows® 2000 operating system:

1. Point-to-Point Tunneling Protocol (PPTP)
   PPTP uses user-level Point-to-Point Protocol (PPP) authentication methods and Microsoft Point-to-Point Encryption (MPPE) for data encryption.
2. Layer Two Tunneling Protocol (L2TP) with Internet Protocol security (IPSec)
   L2TP uses user-level PPP authentication methods and IPSec for computer-level authentication using certificates and data authentication, integrity, and encryption.

A remote access client (a single user computer) makes a remote access VPN connection that connects to a private network. The VPN server provides access to the entire network to which the VPN server is attached. The packets sent from the remote client across the VPN connection originate at the remote access client computer.

The remote access client (the VPN client) authenticates itself to the remote access server (the VPN server) and, for mutual authentication, the server authenticates itself to the client.

Computers running Windows XP, Windows 2000, Windows NT® version 4.0, Windows Millennium Edition (ME), Windows 98, and Windows 95 operating systems can create remote access VPN connections to a VPN server running Windows 2000. VPN clients may also be any non-Microsoft PPTP client or L2TP client using IPSec.

Note: Using IPSec tunnel mode is not a remote access VPN technology supported by Microsoft VPN clients or servers due to the lack of an industry standard method of performing user authentication and IP address configuration over an IPSec tunnel. IPSec tunnel mode is described in RFCs 2401, 2402, and 2406.

For encryption, you can use either link encryption or end-to-end encryption in addition to link encryption:

- Link encryption encrypts the data only on the link between the VPN client and the VPN server. For PPTP connections, you must use MPPE in conjunction with MS-CHAP, MS-CHAP v2, or EAP-TLS authentication. For L2TP/IPSec connections, IPSec provides encryption on the link between the VPN client and the VPN server.
- End-to-end encryption encrypts the data between the source host and its final destination. You can use IPSec after the VPN connection is made to encrypt data from the source host to the destination host.

Components of Windows 2000 Remote Access VPNs

Figure 2 shows the components of Windows 2000 remote access virtual private networks.

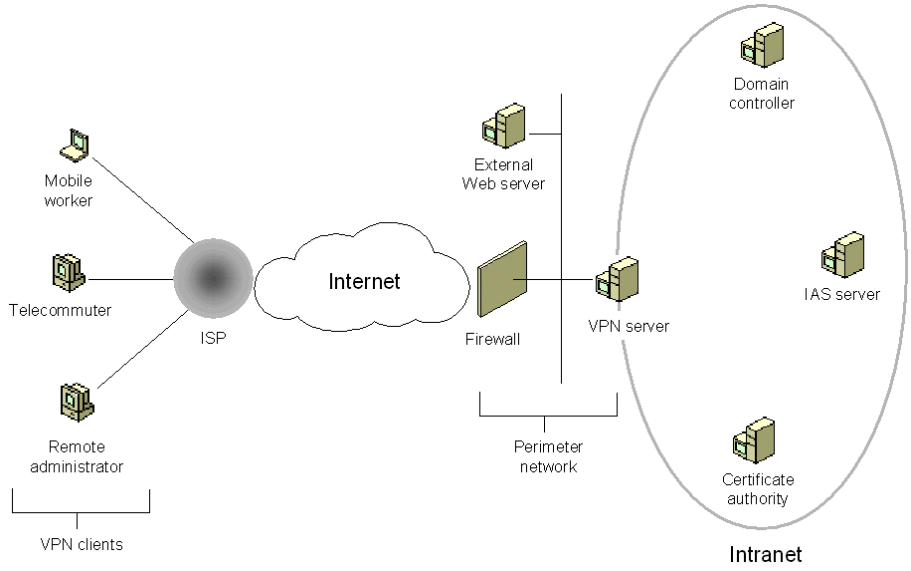

Figure 2 Components of Windows 2000 remote access VPNs

The major components are:

- VPN clients
- Internet infrastructure
- VPN server
- Intranet infrastructure
- Authentication, authorization, and accounting (AAA) infrastructure
- Certificate infrastructure

VPN Clients

The VPN client can be any computer that is capable of creating a PPTP connection using MPPE or L2TP connection using IPSec encryption. Table 1 lists the VPN-capable Microsoft operating systems.

Table 1 VPN-Capable Microsoft Operating Systems

VPN Tunneling Protocol	Microsoft Operating System
PPTP	Windows XP, Windows 2000, Windows NT version 4.0, Windows ME, Windows 98, Windows 95 (with the Windows Dial-Up Networking 1.3 or later Performance & Security Update)

L2TP/IPSec	Windows XP, Windows 2000, and Windows NT 4.0 Workstation, Windows ME, and Windows 98 with Microsoft L2TP/IPSec VPN Client

Typical VPN clients are:

- Laptop users who connect to the organization intranet to access email and other resources while traveling.
- Telecommuters who use the Internet to access organization resources from home.
- Remote administrators who use the Internet to connect to an organization network and configure network or application services.

Microsoft VPN clients can configure VPN connections either manually or by using the Connection Manager components available in Windows 2000. To manually configure a Windows 2000 VPN client, use Make New Connection in the Network and Dial-up Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet. To manually configure a Windows XP VPN client, use the New Connection Wizard in the Network Connections folder to create a VPN connection to the IP address or DNS name of the VPN server on the Internet.

Connection Manager

When scaling the configuration of VPN connections for an enterprise, there are the following problems:

- The exact procedure to configure a VPN connection varies depending on the version of Windows running on the client computer.
- To prevent configuration errors, it is preferable to have the information technology (IT) staff configure the VPN connection rather than end users.
- A configuration method must be able to scale to hundreds or thousands of client computers in a large organization.
- A VPN connection may need a double-dial configuration, where a user must dial the Internet first before creating a VPN connection with the organization intranet.

The solution to these issues of configuring VPN connections across an enterprise is Connection Manager. Connection Manager consists of the following:

- Connection Manager
- Connection Manager Administration Kit
- Connection Point Services

Connection Manager

Connection Manager is a client dialer, included in Windows 2000, whose advanced features make it a superset of basic dial-up networking. Windows 2000 Server includes a set of tools that enables a network manager to deliver pre-configured connections to network users. These tools are the Connection Manager Administration Kit (CMAK) and Connection Point Services (CPS).

Connection Manager provides support for local and remote connections to your service using a network of access points, such as those available worldwide through ISPs. If your service requires secure connections over the Internet, you can also use Connection Manager to establish VPN connections to your service.

Connection Manager Administration Kit

A network administrator can tailor the appearance and behavior of a connection made with Connection Manager by using CMAK. With CMAK, an administrator can develop client dialer and connection software that allows users to connect to the network by using only the connection features that the administrator defines for them. Connection Manager supports a variety of features that both simplify and enhance implementation of connection support for you and your users, most of which can be incorporated using the Connection Manager Administration Kit Wizard.

CMAK allows you to build profiles customizing the Connection Manager installation package that you deliver to your customers, so that Connection Manager reflects the identity of your organization. It allows you to determine which functions and features you want to include and how Connection Manager appears to your customers. You can do this by using the Connection Manager Administration Kit Wizard to build custom service profiles.

For more information about CMAK and the configuration of connection manager service profiles, see Windows 2000 Server Help.

Connection Point Services

Connection Point Services (CPS) enables you to automatically distribute and update custom phone books. These phone books contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an Internet access point. The phone books give users complete POP information, so when they travel they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would have to contact their organization's technical support staff to be informed of changes in POP information and to reconfigure their client dialer software.

CPS has two components:

1. Phone Book Administrator
   A tool used to create and maintain the phone book database and to publish new phone book information to the Phone Book Service.
2. Phone Book Service
   A Microsoft Internet Information Services (IIS) extension that runs on Windows NT Server 4.0 or later (with IIS). Phone Book Service automatically checks subscribers' or corporate employees' current phone books and, if necessary, downloads a phone book update.

For more information about CPS and the configuration of phone books, see Windows 2000 Server Help.

Single sign-on

Single sign-on is the capability that allows a remote access user to create a remote access connection to an organization and logon to the organization's domain by using the same set of credentials. For a domain-based infrastructure, the user name and password or smart card is used for both authenticating and authorizing a remote access connection and for authenticating and logging on to a Windows domain. Single sign-on is performed by selecting the **Logon by using dial-up networking option** on the Windows XP and Windows 2000 logon dialog box and then selecting a dial-up or VPN connection to use to connect to the organization.

For VPN connections, the user must first connect to the Internet before creating a VPN connection. After the Internet connection is made, the VPN connection and logon to the domain can be accomplished. If there is a separate ISP account that the user uses to connect to the Internet, you can create a dial-up connection with the ISP

credentials already configured. Then, configure your VPN connection to dial the ISP connection before attempting the VPN connection. In this configuration, the user will never have to type the ISP credentials when logging on to the domain. This association between the VPN connection and the ISP connection can be configured manually or by using Connection Manager.

Installing a certificate on a client computer

If your Windows 2000 VPN clients are either making L2TP connections or using certificates for user-level authentication, certificates must be installed on the VPN client computer. For L2TP connections, a computer certificate must be installed on the VPN client computer to provide authentication for establishing an IPSec security association (SA). For user-level authentication using the Extensible Authentication Protocol-Transport Level Security (EAP-TLS) authentication protocol, you can either use a user certificate or a smart card.

For user certificate-based authentication, the computer user must request a user certificate from a Windows 2000 certification authority (CA) on your intranet. For smart card-based authentication, a network administrator must configure an enrollment station and issue smart cards with certificates that are mapped to individual user accounts.

For more information about installing certificates on VPN client computers, see “Certificate Infrastructure” in this paper.

Design Points: Configuring the VPN client

Consider the following when configuring your VPN clients for remote access VPN connections:

- If you have a small number of VPN clients, perform manual configuration of VPN connections on each computer.
- If you have a large number of VPN clients running different versions of Microsoft operating systems, use the Connection Manager components of Windows 2000 to create the custom VPN connection configuration package for distribution and to maintain the phone book database for your POPs.
- If you are using Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client to make L2TP connections, you must install a computer certificate on the VPN client computer.
- If you are using Windows XP or Windows 2000 VPN clients and user-level certificate authentication with EAP-TLS, you must either install a user certificate on the VPN client computer or a user certificate on the smart card used by the VPN client computer.

Internet Network Infrastructure

To create a VPN connection to a VPN server across the Internet:

- The VPN server's name must be resolvable.
- The VPN server must be reachable.
- VPN traffic must be allowed to and from the VPN server.

VPN server name resolvability

In most cases you want to reference the VPN server by name, rather than an IP address, as names are much easier to remember. You can use a name (for example VPN1.example.microsoft.com) as long as the name can be resolved to an IP address. Therefore, you must ensure that whatever name you are using for your VPN servers when configuring a VPN connection, that name must be able to be resolved to an IP address using the Internet Domain Name System (DNS) infrastructure.

When you use names rather than addresses, you can also take advantage of DNS round robin load balancing if you have multiple VPN servers with the same name. Within DNS, you can create multiple records that resolve a specific name to different IP address. In this situation, DNS servers send back all the addresses in response to a DNS name query and randomize the order of the addresses for successive queries. Because most DNS clients use the first address in the DNS query response, the result is that VPN client connections are on average spread across the VPN servers.

VPN server reachability

To be reachable, the VPN server must be assigned a public IP address to which packets are forwarded by the routing infrastructure of the Internet. If you have been assigned a static public IP address from an ISP or an Internet registry, this is typically not an issue. In some configurations, the VPN server is actually configured with a private IP address and has a published static IP address by which it is known on the Internet. A device between the Internet and the VPN server translates the published and actual IP addresses of the VPN server in packets to and from the VPN server.

While the routing infrastructure might be in place, the VPN server might be unreachable due to the placement of firewalls, packet filtering routers, network address translators, security gateways, or other types of devices that prevent packets from either being sent to or received from the VPN server computer.

VPN servers and firewall configuration

There are two approaches to using a firewall with a VPN server:

1. The VPN server is attached directly to the Internet and the firewall is between the VPN server and the intranet.
   In this configuration, the VPN server must be configured with packet filters that only allow VPN traffic in and out of its Internet interface. The firewall can be configured to allow specific types of remote access traffic.
2. The firewall is attached to the Internet and the VPN server is between the firewall and the intranet.
   In this configuration, both the firewall and the VPN server are attached to a network segment known as the perimeter network (also known as a demilitarized zone [DMZ] or a screened subnet). Both the firewall and the VPN server must be configured with packet filters that allow only VPN traffic to and from the Internet. Figure 2 shows this configuration.

For the details of configuring packet filters for the VPN server and the firewall for both of these configurations, see Appendix A.

Design Points: VPN server accessibility from the Internet

Consider the following when configuring your Internet infrastructure for remote access VPN connections:

- Ensure that the DNS name of your VPN servers are resolvable from the Internet by either placing an appropriate DNS record in your Internet DNS server or the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your VPN server when directly connected to the Internet. Due to packet filtering, the result of the ping command may be "Request timed out", but check to ensure that the name specified was resolved by the Ping tool to the proper IP address.
- Ensure that the IP addresses of your VPN servers are reachable from the Internet by using the Ping tool to ping the name or address of your VPN server with a 5 second timeout (using the -w command line

option) when directly connected to the Internet. If you see a "Destination unreachable" error message, the VPN server is not reachable.
- Configure packet filtering for PPTP traffic, L2TP traffic, or both types of traffic on the appropriate firewall and VPN server interfaces connecting to the Internet and the perimeter network. For more information, see Appendix A.

Authentication Protocols

To authenticate the user who is attempting a create a PPP connection, Windows 2000 supports a wide variety of PPP authentication protocols including:

- Password Authentication Protocol (PAP)
- Challenge-Handshake Authentication Protocol (CHAP)
- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)
- MS-CHAP version 2 (MS-CHAP v2)
- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)
- Extensible Authentication Protocol-Transport Level Protocol (EAP-TLS)

For PPTP connections, you must use MS-CHAP, MS-CHAP v2, or EAP-TLS. Only these three authentication protocols provide a mechanism to generate the same encryption key on both the VPN client and the VPN server. MPPE uses this encryption key to encrypt all PPTP data sent on the VPN connection. MS-CHAP and MS-CHAP v2 are password-based authentication protocols.

In the absence of user certificates or smart cards, MS-CHAP v2 is highly recommended as it is a stronger authentication protocol than MS-CHAP and provides mutual authentication. With mutual authentication, the VPN client is authenticated by the VPN server and the VPN server is authenticated by the VPN client.

Note: If you must use a password-based authentication protocol, enforce the use of strong passwords on your network. Strong passwords are long (greater than 8 characters) and contain a random mixture of upper and lower case letters, numbers, and punctuation. An example of a strong password is f3L*q02~>xR3w#4o. In an Active Directory™ service domain, use Group Policy settings to create and enforce strong user passwords.

EAP-TLS is designed to be used in conjunction with a certificate infrastructure and either user certificates or smart cards. With EAP-TLS, the VPN client sends its user certificate for authentication and the VPN server sends a computer certificate for authentication. This is the strongest authentication method as it does not rely on passwords.

Note: You can use third-party CAs as long as the certificate in the computer store of the IAS server contains the Server Authentication certificate purpose (also known as a certificate usage or certificate issuance policy). A certificate purpose is identified using an object identifier (OID). The OID for Server Authentication is "1.3.6.1.5.5.7.3.1". Additionally, the user certificate installed on the Windows 2000 remote access client must contain the Client Authentication certificate purpose (OID "1.3.6.1.5.5.7.3.2").

For L2TP/IPSec connections, any authentication protocol can be used because the authentication occurs after the VPN client and VPN server have established a secure channel of communication known as an IPSec security association (SA). However, the use of either MS-CHAP v2 and EAP-TLS are recommended to provide strong user authentication.

Design Point: Which authentication protocol to use?
Consider the following when choosing an authentication protocol for VPN connections:

- If you are using smart cards or have a certificate infrastructure that issues user certificates, use the EAP-TLS authentication protocol for both PPTP and L2TP connections. EAP-TLS is only supported by VPN clients running Windows XP and Windows 2000.
- If you must use a password-based authentication protocol, use MS-CHAP v2 and enforce strong passwords using group policy. MS-CHAP v2 is supported by computers running Windows XP, Windows 2000, Windows NT 4.0 with Service Pack 4 and later, Windows ME, Windows 98, and Windows 95 with the Windows Dial-Up Networking 1.3 or later Performance & Security Update.

VPN Protocols

Windows 2000 includes support for two remote access VPN protocols:

1. Point-to-Point Tunneling Protocol
2. Layer Two Tunneling Protocol

Point-to-Point Tunneling Protocol
Introduced in Windows NT 4.0, PPTP leverages Point-to-Point Protocol (PPP) user authentication and Microsoft Point-to-Point Encryption (MPPE) to encapsulate and encrypt IP, IPX, and NetBEUI traffic. When version 2 of the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2) is used with strong passwords, PPTP is a secure VPN technology. For nonpassword-based authentication, Extensible Authentication Protocol-Transport Level Security (EAP-TLS) can be used in Windows 2000 to support smart cards. PPTP is widely supported, easily deployed, and can be used across network address translators (NATs).

Layer Two Tunneling Protocol with IPSec
L2TP leverages PPP user authentication and IPSec encryption to encapsulate and encrypt IP, IPX, and NetBEUI traffic. This combination, known as L2TP/IPSec, uses certificate-based computer identity authentication to create the IPSec security association in addition to PPP-based user authentication. L2TP/IPSec provides data integrity and data authentication for each packet. However, L2TP/IPSec requires a certificate infrastructure to allocate computer certificates and is supported by Windows XP, Windows 2000, and Microsoft L2TP/IPSec VPN Client L2TP clients.

Design Point: PPTP or L2TP?
Consider the following when deciding between PPTP and L2TP for remote access VPN connections:

- PPTP can be used with a variety of Microsoft clients including Windows XP, Windows 2000, Windows NT version 4.0, Windows ME, Windows 98, and Windows 95 with the Windows Dial-Up Networking 1.3 or later Performance & Security Update. PPTP does not require a certificate infrastructure to issue computer certificates.
- PPTP-based VPN connections provide data confidentiality (captured packets cannot be interpreted without the encryption key). PPTP VPN connections, however, do not provide data integrity (proof that the data was not modified in transit) or data authentication (proof that the data was sent by the authorized user).
- PPTP-based VPN clients can be located behind a NAT if the NAT includes a NAT editor that knows how

to properly translate PPTP tunneled data. For example, the Internet connection sharing (ICS) feature of the Network and Dial-up Connections folder and the NAT routing protocol component of the Routing and Remote Access service include a NAT editor that translates PPTP traffic from PPTP clients located behind the NAT. VPN servers cannot be behind a NAT unless there are multiple public IP addresses and there is a one-to-one mapping of a public IP address to the private IP address of the VPN server or, if there is only one public address, if the NAT is configured to translate and forward the PPTP tunneled data to the VPN server. Most NATs using a single public IP address, including ICS and the NAT routing protocol component, can be configured to allow inbound traffic based on IP addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, a VPN server cannot be located behind a computer using ICS or the NAT routing protocol component when using a single IP address.

- L2TP-based VPN clients or servers cannot be behind a NAT unless both support IPSec NAT Traversal (NAT-T). IPSec NAT-T is supported by Windows Server 2003, Microsoft L2TP/IPSec VPN Client, and for VPN clients with L2TP/IPSec NAT-T Update for Windows XP and Windows 2000. Windows 2000 Server does not support IPSec NAT-T.
- L2TP can be used with Windows XP, Windows 2000, and Microsoft L2TP/IPSec VPN Client L2TP clients and supports computer certificates as the recommended authentication method for IPSec. Computer certificate authentication requires a certificate infrastructure to issue computer certificates to the VPN server computer and all VPN client computers.
- By using IPSec, L2TP-based VPN connections provide data confidentiality, data integrity, data authentication, and replay protection.
- PPTP and L2TP is not an either/or choice. By default, a Windows 2000 VPN server supports both PPTP and L2TP connections simultaneously. You can use PPTP for some remote access VPN connections (from VPN clients that are not running Windows XP or Windows 2000 and do not have an installed computer certificate) and L2TP for other remote access VPN connections (from VPN clients running Windows XP, Windows 2000, or Microsoft L2TP/IPSec VPN Client and have an installed computer certificate).
- If you are using both PPTP and L2TP, you can create separate remote access policies that define different connection parameters for PPTP and L2TP connections.

VPN Server

A VPN server is a computer running Windows 2000 Server and the Routing and Remote Access service. The VPN server does the following:

- Listens for PPTP connection attempts and IPSec SA negotiations for L2TP connection attempts.
- Authenticates and authorizes VPN connections before allowing data to flow.
- Acts as a router forwarding data between VPN clients and resources on the intranet.
- Acts as an endpoint of the VPN tunnel from the tunnel client (typically the VPN client).
- Acts as the endpoint of the VPN connection from the VPN client.

The VPN server typically has two or more installed network adapters—one or more network adapters connected to the Internet and one or more network adapters connected to the intranet. The configuration of a VPN server with a single network adapter is discussed in Appendix B.

When you configure and enable the Routing and Remote Access service, the Routing and Remote Access Server

Setup Wizard prompts you to select the role that the computer will fulfill. For VPN servers, you should select the **Virtual private network (VPN) server** option.

With the **Virtual private network (VPN) server** option, the Routing and Remote Access server operates in the role of a VPN server that supports remote access and router-to-router VPN connections. For remote access VPN connections, users run VPN client software and initiate a remote access connection to the server. For router-to-router VPN connections, a router initiates a VPN connection to the server. Alternately, the server initiates a VPN connection to another router.

When you select the **Virtual private network (VPN) server** option in the Routing and Remote Access Server Setup Wizard:

1. You are first prompted to verify the protocols over which VPN traffic is forwarded. By default, all of the protocols that can be used with a remote access or router-to-router VPN connection are listed.
2. Next, you are prompted to select the interface that is connected to the Internet. If the VPN server is not connected to the Internet, you can select **<No Internet connection>**. The interface that you select will be automatically configured with packet filters that allow only PPTP and L2TP-related traffic. All other traffic is silently discarded. For example, you will no longer be able to ping the Internet interface of the VPN server. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.
3. Next, if you have multiple network adapters that are connected to the intranet, you are prompted to select an interface over which DHCP, DNS, and WINS configuration is obtained.
4. Next, you are prompted to determine whether you want to assign IP addresses to either remote access clients or calling routers by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add one or more address ranges.
5. Next, you are prompted to specify whether you want to use RADIUS as your authentication and accounting provider. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the shared secret.

When you select the **Virtual private network (VPN) server** option in the Routing and Remote Access Server Setup Wizard, the results are as follows:

1. The Routing and Remote Access service is enabled as both a remote access server and a LAN and demand-dial router, with Windows as the authentication and accounting provider (unless RADIUS was chosen and configured). If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the IP interface from which to obtain DHCP, DNS, and WINS configuration. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration. If specified, the static IP address ranges are configured.
2. Exactly 128 PPTP and 128 L2TP ports are created. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections.
3. The selected Internet interface is configured with input and output IP packet filters that allow only PPTP and L2TP traffic.
4. All protocols selected are configured to both allow remote access connections and access the network to which the remote access server is attached.
5. The DHCP Relay Agent component is added with the **Internal** interface. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IP address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IP address of a DHCP server on your intranet. The DHCP Relay Agent forwards DHCPInform packets

between VPN remote access clients and an intranet DHCP server.

6. The IGMP component is added. The **Internal** interface and all other LAN interfaces are configured for IGMP router mode. This allows VPN remote access clients to send and receive IP multicast traffic.

Design Points: Configuring the VPN Server

Consider the following before running the Routing and Remote Access Server Setup Wizard:

- Which protocols will be supported over the VPN connection?
  The Routing and Remote Access service can forward IP, IPX, and NetBEUI packets over a PPTP or L2TP connection.
- Which connection of the VPN server is connected to the Internet?
  Typical Internet-connected VPN servers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network) and one connected to the organization intranet. To make this distinction easier to see during the Routing and Remote Access Server Setup Wizard, rename the connections with their purpose or role using the Network and Dial-up Connections folder. For example, rename the connection connected to the Internet, default name **Local Area Connection 2**, to **Internet**.
- Can the VPN server be a DHCP client?
  The VPN server must have a manual TCP/IP configuration for its Internet interface. While technically possible, it is not recommended that the VPN server be a DHCP client for its intranet interface(s). Due to the routing requirements of the VPN server, manually configure an IP address, subnet mask, DNS server(s), and WINS server(s), but do not configure a default gateway.

  Note that it is possible for the VPN server to have a manual TCP/IP configuration and still use DHCP to obtain IP addresses for VPN clients.
- How will IP addresses be allocated to remote access VPN clients?
  The VPN server can be configured to obtain IP addresses from DHCP or from a manually configured set of address ranges. Using DHCP to obtain IP addresses simplifies the configuration, however, you must ensure that the DHCP scope for the subnet to which the intranet connection of the VPN server is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of PPTP and L2TP ports. For example, if the subnet to which the intranet connection of the VPN server is attached contains 50 DHCP clients, then, for the default configuration of the VPN server, the scope must contain at least 307 addresses (50 computers + 128 PPTP clients + 128 L2TP clients + 1 address for the VPN server). If there are not enough IP addresses in the scope, VPN clients that connect after all the addresses in the scope are allocated will be unable to access intranet resources.

  If you are configuring a static pool of addresses, there might be additional routing considerations. For more information, see “Intranet network infrastructure” in this paper.
- What is the authentication and accounting provider?
  The VPN server can use Windows or RADIUS as its authentication or accounting provider.

  When Windows is used as the authentication and accounting provider, the VPN server uses Windows 2000 mechanisms to validate the credentials of the VPN client and access the VPN client's user account dial-in properties. Locally configured remote access policies authorize the VPN connection and locally written accounting log files log VPN connection accounting information.

  When RADIUS is used as the authentication and accounting provider, the VPN server uses a configured

RADIUS server to validate the credentials of the VPN client, authorize the connection attempt, and store VPN connection accounting information.

- Will there be multiple VPN servers?
  If so, create multiple DNS A records to resolve the same name of the VPN server (for example, vpn.microsoft.com) to the different IP addresses of the separate VPN servers. DNS round robin will distribute the VPN connections across the VPN servers.

Consider the following when changing the default configuration of the VPN server for remote access VPN connections:

- Do you need additional PPTP or L2TP ports?
  By default, the Routing and Remote Access Server Setup Wizard configures 128 PPTP and 128 L2TP ports allowing 128 simultaneous PPTP connections and 128 simultaneous L2TP connections. If this is not sufficient for the maximum number of PPTP or L2TP connections, you can change the number of PPTP and L2TP ports by configuring the **WAN miniport (PPTP)** and **WAN miniport (L2TP)** devices from the properties of the **Ports** object in the Routing and Remote Access snap-in.
- Do you need to install a computer certificate?
  If the VPN server is configured with the Windows authentication provider and is supporting L2TP connections or is authenticating connections using the EAP-TLS authentication protocol, you must install a computer certificate on the VPN server that can be validated by the VPN client and a root certificate that is used to validate the VPN client.
- Do you need custom remote access policies for VPN connections?
  If you configure the VPN server for Windows authentication or for RADIUS authentication and the RADIUS server is a computer running Windows 2000 and the Internet Authentication Service (IAS), the default remote access policy rejects all types of connection attempts unless the remote access permission of the user account's dial-in properties is set to **Allow access**. If you want to manage authorization and connection parameters by group or by type of connection, you must configure custom remote access policies. For more information, see "Remote Access Policies" in this paper.
- Do you want separate authentication and accounting providers?
  The Routing and Remote Access Server Setup Wizard configures both authentication and accounting providers to be the same. After the Wizard is complete, however, you can configure the authentication and accounting providers separately (for example, if you want to use Windows authentication and RADIUS accounting). You can configure authentication and accounting providers on the **Authentication** tab from the properties of the VPN server in the Routing and Remote Access snap-in.

Intranet Network Infrastructure

The network infrastructure of the intranet is an important element of VPN design. Without proper design, VPN clients are unable to obtain proper IP addresses and resolve intranet names, and packets cannot be forwarded between VPN clients and intranet resources.

Name resolution

If you use Domain Name System (DNS) to resolve intranet host names or Windows Internet Name Service (WINS) to resolve intranet NetBIOS names, ensure that the VPN server is configured with the IP addresses of the

appropriate DNS and WINS servers. The VPN server can be configured with DNS and WINS server either manually or as a DHCP client. As part of the PPP negotiation process, the VPN clients receive the IP addresses of DNS and WINS server. By default, the VPN clients inherit the DNS and WINS server addresses configured on the VPN server.

After the PPP connection negotiation is complete, Windows XP and Windows 2000 VPN clients send a DHCPInform message to the VPN server. The response is relayed back to the VPN client and contains a DNS domain name, additional DNS server addresses for DNS servers that are checked before the DNS server configured through the PPP negotiation, and WINS server addresses that replace the WINS server addresses configured through the PPP negotiation. This communication is facilitated by the DHCP Relay Agent routing protocol component of the Routing and Remote Access service, which is automatically added by the Routing and Remote Access Server Setup Wizard.

If the VPN server is a DHCP client (the VPN server is using DHCP to configure its intranet interfaces), the VPN server relays the DHCPInform messages to the DHCP server that was in use when the Routing and Remote Access Server Wizard was run. If the VPN server has a manual TCP/IP configuration on its intranet interface (recommended), the DHCP Relay Agent routing protocol component must be configured with the IP address of at least one DHCP server on your intranet. You can add DHCP server IP addresses to the DHCP Relay Agent routing protocol component on the **General** tab from the properties of the **DHCP Relay Agent** object under **IP Routing** in the Routing and Remote Access snap-in.

Design Points: Name resolution by VPN clients for intranet resources

Consider the following when configuring name resolution for remote access VPN clients:

- Using the Ping and Net tools, test DNS and WINS name resolution for intranet resources from the VPN server computer. If name resolution does not work from the VPN server, it will not work for VPN clients. Troubleshoot and fix all name resolution problems of the VPN server before testing VPN connections.
- If the VPN server is a DHCP client (the VPN server is using DHCP to configure its intranet interfaces), no other configuration is necessary. The DNS and WINS servers assigned to the VPN server are also assigned to the VPN clients. The default configuration of the Routing and Remote Access Server Setup Wizard adds the DHCP Relay Agent routing protocol component and configures it with the IP address of the VPN server's DHCP server so that DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 and its response are properly relayed between the VPN client and the DHCP server of the VPN server.

  However, configuring the VPN server as a DHCP client is not recommended due to issues with configuring the VPN server's default gateway. Therefore, it is recommended that you manually configure the TCP/IP configuration of the VPN server's intranet interfaces and manually configure the DHCP Relay Agent routing protocol component with the IP address of one or more of your DHCP servers.
- If the VPN server is manually configured with a TCP/IP configuration, verify the DNS and WINS server addresses. In this configuration, the Routing and Remote Access Server Setup Wizard cannot automatically configure the DHCP Relay Agent routing protocol component. You must manually add the IP address of at least one DHCP server on your intranet in order for DHCPInform messages to be replayed between VPN clients running Windows XP and Windows 2000 and the DHCP server. If you do not, DHCPInform messages sent by VPN clients running Windows XP and Windows 2000 are discarded and the VPN clients do not receive the updated DNS and WINS server addresses or the DNS domain name.

- If you have a single-subnet small office/home office (SOHO) with no DHCP, DNS, or WINS server, you must either configure a DNS server or WINS server in order to resolve names for both computers on the SOHO subnet and VPN clients or use NetBEUI as a LAN protocol across the remote access VPN connection.

  SOHO clients resolve each other's names using a TCP/IP-based local name query broadcast on the SOHO subnet. VPN clients send the same local name query, however, the name query packet is not also broadcast on the SOHO subnet. The result is that while all SOHO clients can resolve each other's names, VPN client computers cannot resolve names without a name server such as a DNS or WINS server. Alternately, you can manually configure a Hosts or Lmhosts file and distribute the file to all the VPN client computers, however, SOHO subnet computers would be unable to resolve the names of VPN client computers.

  By using the NetBEUI protocol across the VPN connection, name resolutions sent using NetBEUI are received by the VPN server and resent over all NetBIOS-based protocols installed on the VPN server using the NetBIOS gateway component of the Routing and Remote Access service. This includes sending the name request as a TCP/IP local name query broadcast on the local subnet. To install NetBEUI support on the VPN client, install the NetBEUI protocol and ensure that it is enabled on the VPN connection. To install NetBEUI support on the VPN server, install the NetBEUI protocol and enable it for remote access connections on the **NetBEUI** tab from the properties of the VPN server in the Routing and Remote Access snap-in. The NetBEUI protocol is not supported by computers running Windows XP.

Routing

The VPN server is an IP router and as such must be properly configured with the set of routes that makes all locations reachable. Specifically, the VPN server needs the following:

- A default route that points to a firewall or router directly connected to the Internet.
  This route makes all of the locations on the Internet reachable.
- One or more routes that summarize the addresses used on your intranet that points to a neighboring intranet router.
  These routes make all of the locations on your intranet reachable from the VPN server. Without these routes, all intranet hosts not connected to the same subnet as the VPN server are unreachable.

To add a default route that points to the Internet, configure the Internet interface with a default gateway and then manually configure the intranet interface without a default gateway.

To add intranet routes to the routing table of the VPN server, you can:

- Add static routes using the Routing and Remote Access snap-in. You do not necessarily have to add a route for each subnet in your intranet. At a minimum, you just need to add the routes that summarize all the possible addresses in your intranet. For example, if your intranet uses portions of the private address space 10.0.0.0/8 to number its subnets and hosts, you do not have to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the intranet subnet to which your VPN server is attached.
- If you are using the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocols in your intranet, you can add and configure the RIP or OSPF routing protocol components of the Routing and Remote Access service so that the VPN server participates in the propagation of routing information as a dynamic router.

If your intranet has only a single subnet, no further configuration is required.

Ensuring the reachability of VPN clients from the intranet depends on how you configure the VPN server to obtain IP addresses for VPN clients. The IP addresses assigned to VPN clients as they connect can be from:

- An on-subnet address range, which is an address range of the intranet subnet to which the VPN server is attached.
  An on-subnet address range is used whenever the VPN server is configured to use DHCP to obtain IP addresses for VPN clients and when the manually configured pool(s) of IP addresses are within the range of addresses of the attached subnet.
- An off-subnet address range, which is an address range that represents a different subnet that is logically attached to the VPN server.
  An off-subnet address range is used whenever the VPN server is manually configured with pool(s) of IP addresses for a separate subnet.

If you are using an on-subnet address range, no additional routing configuration is required as the VPN server acts as a proxy for all packets destined for VPN clients. Routers and hosts on the VPN server subnet forward packets destined to VPN clients to the VPN server and the VPN server relays them the appropriate VPN client.

If you are using an off-subnet address range, you must add the route(s) that summarize the off-subnet address range to the intranet routing infrastructure so that traffic destined to VPN clients are forwarded to the VPN server and then sent by the VPN server to the appropriate VPN client. You can add the routes that summarize the off-subnet address range to the routing infrastructure of the intranet through the following:

- Add static routes to the neighboring router for the off-subnet address range that point to the VPN server's intranet interface. Configure the neighboring router to propagate this static route to other routers in the intranet using the dynamic routing protocol used in your intranet.
- If the VPN server is using OSPF and participating as a dynamic router, the VPN server must be configured as an autonomous system boundary router (ASBR) so that the static routes of the off-subnet address range are propagated to the other OSPF routers in the intranet.

If your intranet consists of a single subnet, then you must either configure each intranet host for persistent route(s) of the off-subnet address range that point to the VPN server's intranet interface or configure each intranet host with the VPN server as its default gateway. Therefore, it is recommended that you use an on-subnet address pool for a SOHO network consisting of a single subnet.

Routing and multi-use VPN servers

If you want to access services running on the VPN server from VPN clients, whether or not the traffic to those services is sent on the Internet in an encrypted or clear text form depends on which address the VPN client is using to access the VPN server service:

- If the VPN client accesses the service running on the VPN server using an intranet IP address of the VPN server, all traffic is sent encrypted inside the tunnel of the VPN connection.
- If the VPN client accesses the service running on the VPN server using the public IP address of the VPN server, all traffic is sent as clear text outside the tunnel of the VPN connection.

Due to the way in which routes are created on VPN remote access clients when making a VPN connection, it may be possible to connect to services running on a VPN server, however, the traffic might not be sent across the VPN connection. When a remote access VPN client creates a VPN connection with a VPN server, it creates the following

routes in the VPN client's IP routing table:

- Default route that uses the VPN connection
  The new default route for the VPN connection effectively replaces the existing default route for the duration of the connection. After the connection is made, all traffic that does not match an address on the directly connected network or the address of the VPN server is sent over the VPN connection.
- Host route to the VPN server that uses the local network connection
  The host route for the VPN server's address is created so that the VPN server is reachable using the locally attached network. If the host route is not present, VPN traffic to the VPN server cannot be sent.

The result of having the host route for the VPN server is that all traffic that is sent between applications running on the VPN client and applications running on the VPN server using the VPN server's public IP address are not sent across the VPN connection, and are instead sent in an unencrypted form across the network between the VPN client and VPN server.

For example, when a remote access VPN client creates a VPN connection with a VPN server and then accesses a file share on the VPN server computer using the VPN server's public IP address, the file sharing traffic is not sent using the VPN connection, but is sent in clear text over the network between the VPN client and VPN server.

Additionally, if packet filters are configured on the VPN server that only allow VPN connection traffic, all other traffic sent to the VPN server is discarded. In this typical configuration, all attempts to connect to services running on the VPN server will fail because traffic attempting to connect to those services are not sent over the VPN connection.

The key to which address is used by the VPN client to access services running on the VPN server lies in the way that the name of the VPN server is resolved. Typical users and applications refer to network resources using names, rather than IP addresses. The name must be resolved to an IP address using either DNS or WINS. If the intranet DNS and WINS infrastructures never contain a record mapping the VPN server's name to the VPN server's public IP address, traffic to services running on the VPN server is always sent across the VPN connection.

To prevent the VPN server from dynamically registering the public IP address of its Internet interface in the intranet DNS, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the Dial-up and Network Connections folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **DNS** tab, and then clear the **Register this connection's addresses in DNS** check box.

To prevent the VPN server from registering the public IP address of its Internet interface with intranet WINS servers, obtain properties of the **Internet Protocol (TCP/IP)** component of the Internet connection in the Dial-up and Network Connections folder. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **WINS** tab, and then click **Disable NetBIOS over TCP/IP**.

Before the VPN connection is made, the VPN client uses the Internet DNS infrastructure to resolve the name of the VPN server computer to its public IP addresses. After the VPN connection is made, assuming that intranet DNS and WINS servers are configured either during the PPP connection process or through the relaying of the DHCPInform message, the VPN client uses the intranet DNS and WINS infrastructures to resolve the name of the VPN server computer to its intranet IP addresses.

VPN client routing and simultaneous intranet and Internet access

By default, when a Windows-based VPN client makes a VPN connection, it automatically adds a new default route for the VPN connection and modifies the existing default route to have a higher metric. Adding the new default route means that all Internet locations except the IP address of the tunnel server and locations based on other routes are

not reachable for the duration of the VPN connection.

To prevent the default route from being created, obtain properties of the **Internet Protocol (TCP/IP)** component of the VPN connection. Click **Advanced**. In the **Advanced TCP/IP Settings** dialog box, click the **General** tab, and then clear the **Use default gateway on remote network** check box. When the **Use default gateway on remote network** check box is cleared, a default route is not created, however, a route corresponding to the Internet address class of the assigned IP address is created. For example, if the address assigned during the connection process is 10.0.12.119, the Windows 2000 and Windows XP VPN client creates a route for the class-based network ID 10.0.0.0 with the subnet mask 255.0.0.0.

Based on the **Use default gateway on remote network** setting, one of the following occurs when the VPN connection is active:

- Internet locations are reachable and intranet locations are not reachable except those matching the address class of the assigned IP address (the **Use default gateway on remote network** check box is cleared).
- All intranet locations are reachable and Internet locations are not reachable except the address of the VPN server and location available through other routes (the **Use default gateway on remote network** check box is selected).

For most Internet-connected VPN clients, this behavior does not represent a problem because they are typically engaged in either intranet or Internet communication, not both.

For VPN clients who want concurrent access to intranet and Internet resources when the VPN connection is active, you can do one of the following:

- Select the **Use default gateway on remote network** check box (the default setting) and allow Internet access through the organization intranet. Internet traffic between the VPN client and Internet hosts would pass though firewalls or proxy servers as if the VPN client is physically connected to the organization intranet. While there is an impact on performance, this method allows Internet access to be filtered and monitored according to the organization's network policies while the VPN client is connected to the organization network.
- If the addressing within your intranet is based on a single class-based network ID, clear the **Use default gateway on remote network** check box. The best example is when your intranet is using the private IP address space 10.0.0.0/8.
- If the addressing within your intranet is not based on a single class-based network ID, clear the **Use default gateway on remote network** check box and use a command file (.CMD) on the VPN client with route commands to add static routes for the network IDs of your intranet using your assigned IP address as the gateway IP address after the connection is made.

  You can determine your assigned IP address from the display of the Ipconfig command or by double-clicking the VPN connection in the Dial-up and Network Connections folder when the VPN connection is active. In the resulting **Status** dialog box, click the **Details** tab. The VPN client's assigned IP address is listed as **Client IP address**.

Design Points: Routing infrastructure

Consider the following when configuring the routing infrastructure for remote access VPN connections:

- Configure the Internet interface of the VPN server with a default gateway. Do not configure the intranet

interface of the VPN server with a default gateway.

- Add static IP routes to the VPN server that summarize the addresses used in your intranet. Alternately, if you use either RIP or OSPF for your dynamic routing protocol, configure and enable RIP or OSPF on the VPN server. If you use a routing protocol other than RIP or OSPF, such as Interior Gateway Routing Protocol (IGRP), you might configure the VPN server's neighboring intranet router for RIP or OSPF on the interface connected to subnet to which the VPN server is attached and IGRP on all other interfaces.
- Do not install services on your VPN server that you want to access from the Internet. Traffic to these services are sent in plaintext across the Internet and are dropped by the VPN server due to VPN packet filters configured at the VPN server. Instead, install the services you want to access on another computer that is available across the VPN server.
- Configure the VPN server with an on-subnet address range by obtaining IP addresses through DHCP or by manually configuring on-subnet address pools.

AAA Infrastructure

The authentication, authorization, and accounting (AAA) infrastructure exists to:

- Authenticate the credentials of VPN clients.
- Authorize the VPN connection.
- Record the VPN connection creation and termination for accounting purposes.

The AAA infrastructure consists of:

- The VPN server computer
- A RADIUS server computer
- A domain controller

As previously discussed, a Windows 2000 VPN server can be configured with either Windows or RADIUS as its authentication or accounting provider. RADIUS provides a centralized AAA service when you have multiple VPN servers or a mix of heterogeneous dial-up and VPN equipment.

When you configure Windows as the authentication provider, the VPN server performs the authentication of the VPN connection by communicating with a domain controller using a secure remote procedure call (RPC) channel and authorization of the connection attempt through the dial-in properties of the user account and locally configured remote access policies.

When you configure RADIUS as the authentication provider, the VPN server relies on a RADIUS server to perform both the authentication and authorization. When a VPN connection is attempted, the VPN client credentials and other connection parameters are used to create a RADIUS Access-Request message that is sent to the configured RADIUS server. If the connection attempt is both authenticated and authorized, the RADIUS server sends back a RADIUS Access-Accept message. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends back a RADIUS Access-Reject message.

When you configure Windows as the accounting provider, the VPN server logs VPN connection information in a local log file (*SystemRoot*\System32\Logfiles\Logfile.log by default) based on settings configured on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in.

When you configure RADIUS as the authentication provider, the VPN server sends RADIUS accounting messages for VPN connections on a RADIUS server, which records the accounting information.

If you are using RADIUS and a Windows domain as the user account database for which to verify user credentials

and obtain dial-in properties, it is recommended to use the Windows 2000 Internet Authentication Service (IAS). IAS is a full-featured RADIUS server that is tightly integrated with Windows 2000, Active Directory, and the Routing and Remote Access service.

When IAS is used as the RADIUS server:

- IAS performs the authentication of the VPN connection by communicating with a domain controller using a secure RPC channel. IAS performs authorization of the connection attempt through the dial-in properties of the user account and remote access policies configured on the IAS server.
- IAS logs all RADIUS accounting information in a local log file (*SystemRoot*\System32\Logfiles\Logfile.log by default) based on settings configured on the properties of the **Local File** object in the **Remote Access Logging** folder in the Internet Authentication Service snap-in.

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 policies consist of the following elements:

- Conditions
- Permission
- Profile settings

Conditions

Remote access policy conditions are one or more attributes that are compared to the settings of the connection attempt. If there are multiple conditions, then all of the conditions must match the settings of the connection attempt in order for it to match the policy. For VPN connections, you commonly use the following conditions:

- NAS-Port-Type
  By setting the NAS-Port-Type condition to Virtual (VPN), you can specify all VPN connections.
- Tunnel-Type
  By setting the Tunnel-Type to Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP), you can specify different policies for PPTP and L2TP connections.
- Windows-Groups
  By setting the Windows-Groups to the appropriate groups, you can grant or deny access based on group membership.

Permission

You can use the permission setting to either grant to deny remote access for the connection attempt if the remote access permission of the user account is set to **Control access through Remote Access Policy**. Otherwise, the remote access permission setting on the user account determines the remote access permission.

Profile Settings

A remote access policy profile is a set of properties that are applied to a connection when it is authorized. For VPN connections, you can use the following profile settings:

- Dial-in constraints can be used to define how long the connection can exist or be idle before being terminated by the VPN server, among others.
- IP settings can define using IP packet filters the specific types of IP traffic that are allowed for remote access VPN connections. With profile packet filters, you can configure the IP traffic that is allowed from remote access clients (From client filters) or to remote access clients (To client filters) on an exception basis: either all traffic except traffic specified by filters or no traffic except traffic specified by filters. Remote access policy profile filtering applies to all remote access connections that match the remote access policy.
- Authentication settings can define which authentication protocols the VPN client must use in order to send its credentials and the configuration of EAP types, such as EAP-TLS.
- Encryption settings can define whether encryption is required and the encryption strength. For encryption strengths, Windows 2000 supports **Basic** (40-bit MPPE for PPTP and 56-bit Data Encryption Standard [DES] for L2TP), **Strong** (56-bit MPPE for PPTP and 56-bit DES for L2TP), or **Strongest** (128-bit MPPE for PPTP and 3DES for L2TP). **Strongest** can be used only if the Windows 2000 High Encryption Pack or Service Pack 2 and later is installed.

For example, you can create a Windows 2000 group called VPNUsers whose members are the user accounts of the users creating remote access VPN connections across the Internet. Then, you create a policy with two conditions on the policy: NAS-Port-Type is set to Virtual (VPN) and Windows-Group is set to VPNUsers. Finally, you configure the profile for the policy to select a specific authentication method and encryption strength.

Preventing traffic routed from VPN clients

Once a VPN client successfully establishes a PPTP or L2TP connection, by default any packet sent over the connection is received by the VPN server and forwarded. Packets sent over the connection can include:

- Packets originated from the remote access client computer
- Packets sent to the remote access client computer by other computers

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 computer that is forwarding traffic to the remote access VPN client has not been authenticated by the VPN server. 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 sending 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.

For the **From client filter**, set the filter action to **Deny all traffic except those listed below** and configure a single filter with the settings listed in Table 2.

Table 2 From client filter settings

IP Packet Filter Field	Setting
Source Address	User's address
Source Network Mask	User's mask
Destination Address	Any
Destination Network Mask	Any
Protocol	Any

For the **To client filter**, set the filter action to **Deny all traffic except those listed below** and configure a single filter with the settings listed in Table 3.

Table 3 To client filter settings

IP Packet Filter Field	Setting
Source Address	Any
Source Network Mask	Any
Destination Address	User's address
Destination Network Mask	User's mask
Protocol	Any

Note: Although the Routing and Remote Access snap-in displays **User's address** and **User's mask**, the actual filter that is created for each remote access client is for the client's assigned IP address and a subnet mask of 255.255.255.255

With this set of IP packet filters, the VPN server discards all traffic sent across the VPN connection except traffic that either originated from or is sent to authenticated remote access VPN clients.

Windows domain user accounts and groups

Windows NT version 4.0 domains, Windows 2000 mixed-mode Active Directory domains, and Windows 2000 native-mode domains contain the user accounts and groups used by the Routing and Remote Access service and IAS to authenticate and authorize VPN connection attempts.

User accounts contain the user name and a form of the user’s password that can be used for validation of the VPN client’s user credentials. Additional account properties determine whether the user account is enabled or disabled, locked out, or permitted to logon only during specific hours. If a user account is disabled, locked out, or not permitted to logon during the time of the VPN connection, the VPN connection attempt is rejected.

User accounts also contain dial-in settings. The dial-in setting most relevant for VPN connections is the remote access permission setting, which has the following values:

- Allow access
- Deny access

- Control access through Remote Access Policy

The **Allow access** and **Deny access** settings explicitly allow or deny remote access and are equivalent to the remote access permission setting of Windows NT 4.0 domain accounts. When you use the **Control access through Remote Access Policy** setting, the remote access permission is determined by the remote access permission setting of the matching remote access policy. If the user account is in a mixed-mode domain, the **Control access through Remote Access Policy** setting is not available and you must manage remote access permission on a per-user basis. If the user account is in a native-mode domain, the **Control access through Remote Access Policy** setting is available and you can manage remote access permission on a per-user basis or using groups.

When using groups to manage access, you can use your existing groups and create remote access policies that either allow or reject access or restrict access based on the group name. For example, the Employees group has no VPN remote access restrictions, however, the Contractors group can only create VPN connections during business hours. Alternately, you can create groups based on the type of connection being made. For example, you can create a VPNUsers group and add as members all the user accounts allowed to create VPN connections.

Both the Routing and Remote Access service and IAS can use Active Directory universal principal names (UPNs) and universal groups. In a large domain with thousands of users, create a universal group for all of the users for whom you want to allow access, and then create a remote access policy that grants access for this universal group. Do not put all of your user accounts directly into the universal group, especially if you have a large number of them on your network. Instead, create separate global groups that are members of the universal group, and add users to those global groups.

Design Points: AAA infrastructure

Consider the following when configuring the AAA infrastructure for remote access VPN connections:

- If you have multiple VPN servers and you want to centralize AAA service or a heterogeneous mixture of dial-up and VPN equipment, use a RADIUS server and configure the VPN server for the RADIUS authentication and accounting providers.
- If your user account database is a Windows domain, use IAS as your RADIUS server. If you use IAS, install IAS on a domain controller for best performance. Install at least two IAS servers for fail-over and fault tolerance of AAA services.
- Whether configured locally or on an IAS server, use remote access policies to authorize VPN connections and specify connection constraints. For example, use the remote access policy profile settings to grant access based on group membership, to enforce the use of encryption and a specific encryption strength, to specify the use of EAP-TLS, or to limit traffic using IP packet filtering.
- To prevent VPN clients from forwarding routed traffic, configure remote access policy profile packet filters to discard all traffic on VPN connections except traffic to and from VPN clients.
- For a large Active Directory domain, nest global groups within universal groups to manage access based on group membership.
- Sensitive fields of RADIUS messages, such as the user password and encryption keys, are encrypted with the RADIUS shared secret configured on the VPN server and the RADIUS server. Make the shared secret a long (22 characters or longer), random sequence of letters, numbers, and punctuation. An example of a strong shared secret is 8d#>9fq4bV)H7%a3^jfDe2. To further protect RADIUS traffic, use Windows 2000 IPSec policies to provide data confidentiality for all traffic using the RADIUS UDP ports (1812 and 1645 for RADIUS authentication traffic and 1813 and 1646 for RADIUS accounting traffic) .

Certificate Infrastructure

To perform certificate-based authentication for L2TP connections and smart card or user certificate-based authentication for VPN connections using EAP-TLS, a certificate infrastructure, also known as a public key infrastructure (PKI), must be in place to issue the proper certificates to submit during the authentication process and to validate the certificate being submitted.

Computer certificates for L2TP/IPSec

When you are using the certificate authentication method for L2TP connections, the list of certification authorities (CAs) is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPSec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPSec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPSec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPSec security negotiation fails.

The VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client was issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during IPSec security negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPSec security negotiation fails.

A single CA commonly issues computer certificates to all computers in an organization. Because of this, all computers within the organization both have computer certificates from a single CA and request certificates for authentication from the same single CA.

Deploying computer certificates in your organization consists of the following:

1. Deploy a certificate infrastructure. For more information, see Appendix D, "Deploying a Certificate Infrastructure".
2. Install a computer certificate on each computer. For more information, see "Deploying L2TP-based Remote Access" in this paper.

Certificate infrastructure for smart cards

The use of smart cards for user authentication is the strongest form of user authentication in Windows 2000. For remote access VPN connections, you must use the Extensible Authentication Protocol (EAP) with the **Smart card or other certificate (TLS)** EAP type, also known as EAP-Transport Level Security (EAP-TLS).

Deploying smart cards in your organization consists of the following:

1. Create a certificate infrastructure using certification authorities.
2. For each domain, set security permissions and delegation for the Smart Card User, Smart Card Logon and Enrollment Agent certificate templates.
3. Configure the CA to issue smart card and Enrollment Agent certificates.
4. Configure an enrollment station, a computer that is used to physically install the smart card certificates on

smart cards.

5. Use the enrollment station to create a smart card with a smart card user logon certificate that is installed on the smart card and assigned to a specific user account.

For more information on how to configure smart cards for user logon, see the topic titled “Checklist: Deploying smart cards for logging on to Windows” in Windows 2000 Server Help.

To find a specific topic in Windows 2000 Server Help

1. From the Windows 2000 desktop, click **Start**, and then click **Help**.
2. In the **Windows 2000** dialog box, click the **Search** tab.
3. Clear the **Match similar words** check box and select the **Search titles only** check box.
4. In **Type the keyword to find**, type the topic title, and then click **List topics**.
5. In the list of topics under **Select topic**, double click the topic that exactly matches the typed topic title.

The individual smart cards are distributed to users who have a computer with a smart card reader. To log in to the computer, the smart card must be inserted into the smart card reader and the smart card personal identification number (PIN) must be typed. When the user attempts a VPN connection, the smart card certificate is sent during the connection negotiation process.

To configure EAP-TLS for smart cards on the VPN client:

- The VPN connection must be configured to use EAP with the **Smart Card or other certificate** EAP type.
- In the properties of the **Smart Card or other certificate** EAP type, select **Use my smart card**. If you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to ensure that the server’s DNS name ends in a specific string, select **Connect only if server name ends with** and type the string. To require the server’s computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in **Trusted root certificate authority**.

To configure EAP-TLS authentication on the VPN server:

- EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the VPN server in the Routing and Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy:

- On the remote access policy that is being used for VPN connections, EAP must be enabled with the **Smart Card or other certificate** EAP type selected on the **Authentication** tab on the policy’s profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

Certificate infrastructure for user certificates

The use of registry-based user certificates for user authentication can be used in place of smart cards. However, it is not as strong a form of authentication. With smart cards, the user certificate issued during the authentication process is only made available when the user physically possesses the smart card and has knowledge of the PIN to logon to their computer. With user certificates, the user certificate issued during the authentication process is made available when the user logs on to their computer using a domain-based user name and password.

Just as with smart cards, authentication using user certificates for remote access VPN connections use EAP-TLS as the authentication protocol.

Deploying user certificates in your organization consists of the following:

3. Deploy a certificate infrastructure. For more information, see Appendix D, "Deploying a Certificate Infrastructure".
4. Install a user certificate for each user. For more information, see "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" in this paper.

When the user attempts a VPN connection, the user certificate is sent during the connection negotiation process.

To configure EAP-TLS for user certificates on the VPN client:

- The VPN connection must be configured to use EAP with the **Smart Card or other certificate** EAP type.
- In the properties of the **Smart Card or other certificate** EAP type, select **Use a certificate on this computer**. If you want to validate the computer certificate of the VPN or IAS server, select **Validate server certificate**. If you want to ensure that the server's DNS name ends in a specific string, select **Connect only if server name ends with** and type the string. To require the server's computer certificate to have been issued a certificate from a specific trusted root CA, select the CA in **Trusted root certificate authority**.

To configure EAP-TLS authentication on the VPN server:

- EAP must be enabled as an authentication type on the **Authentication Methods** dialog box available from the **Security** tab in the properties of the VPN server in the Routing and Remote Access snap-in.

To configure EAP-TLS authentication on the remote access policy:

- On the remote access policy that is being used for VPN connections, EAP must be enabled with the **Smart Card or other certificate** EAP type selected on the Authentication tab on the policy's profile settings. If the computer on which the remote access policy is being configured has multiple computer certificates installed, configure the properties of the **Smart Card or other certificate** EAP type and select the appropriate computer certificate to submit during the EAP-TLS authentication process.

Design Points: Certificate infrastructure

Consider the following when configuring the certificate infrastructure for remote access VPN connections:

- In order to create L2TP/IPSec remote access VPN connections using computer certificate authentication for IPSec, you must install computer certificates, also known as machine certificates, on each VPN client and VPN server. If you are using a Windows 2000 enterprise CA as an issuing CA, configure your Active Directory domain for auto-enrollment of computer certificates using Computer Configuration group policy. Each computer that is a member of the domain automatically requests a computer certificate when the Computer Configuration group policy is updated.
  The computer certificate of the VPN client must be valid and verifiable by the VPN server—the VPN server must have a root CA certificate for the CA that issued the computer certificate of the VPN client.

  The computer certificate of the VPN server must be valid and verifiable by the VPN client—the VPN client must have a root CA certificate for the CA that issued the computer certificate of the VPN server.
- In order to authenticate VPN connections using a smart card or user certificate with EAP-TLS, the VPN client must have a smart card or registry-based user certificate installed and either the VPN server (if configured for Windows authentication) or the IAS server (if the VPN server is configured for RADIUS authentication and the RADIUS server is a computer running Windows 2000 and IAS) must have a

computer certificate installed.

The smart card or user certificate of the VPN client must be valid and verifiable by the VPN server—the VPN server trust the root CA for the CA that issued the certificate of the VPN client.

The computer certificate of the VPN server must be verifiable by the VPN client—the VPN client trust the root CA for the CA that issued the computer certificate of the VPN server.

- To install a computer or user certificate on a computer across the Internet, make a PPTP connection using a password-based authentication protocol such as MS-CHAP v2. After connecting, use the Certificate Manager snap-in or Internet Explorer to request the appropriate certificates. Once the certificates are installed, disconnect and then reconnect with the appropriate VPN protocol and authentication method. An example of this situation is a laptop computer that is issued to an employee without the certificates needed to make L2TP/IPSec or EAP-TLS-authenticated connections.

Deploying PPTP-based Remote Access

Deploying PPTP-based remote access VPN connections using Windows 2000 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy AAA infrastructure
- Deploy VPN servers
- Deploy intranet infrastructure
- Deploy VPN clients

Deploying Certificate Infrastructure

For PPTP-based VPN connections, a certificate infrastructure is needed only when you are using either smart cards or registry-based user certificates and EAP-TLS authentication. If you are only using a password-based authentication protocol such as MS-CHAP v2, a certificate infrastructure is not required and is not used for the creation of the VPN connection.

If you need a certificate infrastructure for PPTP-based VPN connections, you must install a computer certificate on the authenticating server (the VPN server or the RADIUS server) and either a certificate on each smart card distributed to VPN client users or a user certificate on each VPN client computer.

For information about deploying a certificate infrastructure, see Appendix E, "Deploying a Certificate Infrastructure."

Installing computer certificates

To install a computer certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a computer certificate in the following ways:

1. By configuring the automatic allocation of computer certificates to computers in a Windows 2000 domain.
   This method allows a single point of configuration for the entire domain. All members of the domain automatically request the computer certificate through a group policy setting. To immediately obtain a computer certificate for the VPN or IAS server that is a member of the domain for which auto-enrollment is configured, restart the computer or type **secedit /refreshpolicy machine_policy** from a command prompt. To configure a Windows 2000 domain for automatic enrollment of computer certificates, see the topic titled "Configure automatic certificate allocation from an enterprise CA" in Windows 2000 Server Help.
2. By using the Certificates snap-in to request a computer certificate.
   If you are using a Windows 2000 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA using the Certificates snap-in. For more information, see the topics titled "Manage certificates for a computer" and "Request a certificate" in Windows 2000 Server Help.
3. By using the Certificates snap-in to import a computer certificate.
   If you have a certificate file that contains the computer certificate, you can import the computer certificate using the Certificates snap-in. For more information about importing a certificate using the Certificates snap-in, see the topic titled “Import a certificate” in Windows 2000 Server Help.
4. By executing a CAPICOM script that requests a computer certificate.
   In this method, each computer that needs a computer certificate must execute a CAPICOM script that requests a

computer certificate from the issuing CA. CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft ActiveX® and COM objects. CAPICOM can be used via Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, see CAPICOM at http://msdn.microsoft.com/library/en-us/security/security/capicom_start_page.asp?frame=true.

Deploying smart cards

For information about deploying smart cards in Windows 2000, see the topic titled "Checklist: Deploying smart cards for logging on to Windows" in Windows 2000 Server Help.

Installing user certificates

To install a user certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a user certificate in the following ways:

1. By using a Web browser to request a user certificate.
   The issuing CA must support Web enrollment of certificates. For example, if you are using a Windows 2000 enterprise CA as an issuing CA and the CA computer is also running Internet Information Services (IIS), you can use Web enrollment to request a user certificate. For more information about requesting a user certificate, see the topic titled “Submit a user certificate request via the Web” in Windows 2000 Server Help.
2. By using the Certificates snap-in to request a user certificate.
   If you are using a Windows 2000 enterprise CA as an issuing CA, you can request a user certificate from the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled “Request a certificate” in Windows 2000 Server Help.
3. By importing a user certificate using the Certificates snap-in.
   If you have a certificate file that contains a user certificate, import the user certificate from the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled “Import a certificate” in Windows 2000 Server Help.
4. By executing a CAPICOM script that requests a user certificate.
   In this method, each user must execute a CAPICOM script that requests a user certificate from the issuing CA.

Deploying Internet Infrastructure

Deploying the Internet infrastructure for remote access VPN connections consists of the following:

- Place VPN servers in perimeter network or on the Internet.
- Install Windows 2000 Server on VPN servers and configure Internet interfaces.
- Add address records to Internet DNS.

Placing VPN servers in perimeter network or on the Internet

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between your intranet and the Internet. If so, configure packet filters on the firewall to allow PPTP traffic to and from the IP address of the VPN servers' perimeter network interfaces. For more information, see Appendix A.

Installing Windows 2000 Server on VPN servers and configuring Internet interfaces

Install Windows 2000 Server on the VPN server computer and connect it to either the Internet or to perimeter network with one network adapter and connect it to the intranet with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN server computer will not forward IP packets between the Internet and the intranet. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the Internet.) Do not configure the connection with DNS server or WINS server IP addresses.

Adding address records to Internet DNS

To ensure that the name of the VPN server (for example, vpn.microsoft.com) can be resolved to its proper IP address, either add DNS address (A) records to your DNS server (if you are providing DNS name resolution for Internet users) or have your ISP add DNS address (A) records to their DNS server(s) (if your ISP is providing DNS name resolution for Internet users). Verify that the name of the VPN server can be resolved to its public Internet IP address when connected to the Internet.

Deploying AAA Infrastructure

Deploying the AAA infrastructure for remote access VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.

Configuring Active Directory for user accounts and groups

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all users that are making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to **Allow access** or **Deny access** to manage remote access by user. Or, to manage remote access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**.
3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see the topic titled "Universal, global, and domain local groups" in Windows 2000 Server Help.

Configuring the primary IAS server on a domain controller

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows 2000 Server Help.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next,

configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help. For more information about trust relationships, see the topic titled "Understanding domain trusts" in Windows 2000 Server Help.
If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file properties" in Windows 2000 Server Help.
5. Add the VPN server(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows 2000 Server Help. For the IP address of each VPN server, use the intranet IP address assigned to the VPN server. If you are using names, use the internal name of the VPN server (this is not necessarily the same DNS name used by Internet clients). Use strong shared secrets.
6. Create remote access policies that reflect your remote access usage scenarios.
   For example, to configure a remote access policy that requires PPTP-based VPN connections for members of the Employees group to use EAP-TLS authentication and 128-bit encryption, create a remote access policy with the following settings:

   Policy name: VPN connections

   Conditions:

   NAS-Port-Type matches **Virtual (VPN)**

   Tunnel-Type matches **Point-to-Point Tunneling Protocol**

   Windows-Groups matches **Employees** (example)

   Permission: **Grant remote access permission**

   Profile settings, **Authentication** tab:

   Select **Extensible Authentication Protocol** and the **Smart Card or other Certificate** EAP type. Clear all other check boxes.

   Profile settings, **Encryption** tab:

   Select the **Strongest** check box, and then clear all other check boxes.

7. If you have created new remote access policies, either delete the default remote access policy named **Allow access if dial-up permission is enabled**, or move it so that it is the last policy to be evaluated. For more information, see the topics titled "Delete a remote access policy" and "Change the policy evaluation order" in Windows 2000 Server Help.

Configuring the secondary IAS server on a different domain controller

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows 2000 Server Help.
2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user

objects in Active Directory" in Windows 2000 Server Help.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help. For more information about trust relationships, see the topic titled "Understanding domain trusts" in Windows 2000 Server Help.
   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.
4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see the topic titled "Copy the IAS configuration to another server" in Windows 2000 Server Help.

Deploying VPN Servers

Deploying the VPN servers for remote access VPN connections consists of the following:

- Configure the VPN server's connection to the intranet.
- Run the Routing and Remote Access Server Setup Wizard.

Configuring the VPN server's connection to the intranet

For each VPN server, configure the connection connected to the intranet with a manual TCP/IP configuration consisting of IP address, subnet mask, intranet DNS servers, and intranet WINS servers. Note that you must not configure the default gateway on the intranet connection to prevent default route conflicts with the default route pointing to the Internet.

Running the Routing and Remote Access Server Setup Wizard

Run the Routing and Remote Access Server Setup Wizard to configure each Windows 2000 VPN server using the following steps:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.
2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**.
3. In **Common Configurations**, click **Virtual Private Network (VPN) server** and then click **Next**. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.
4. In **Remote Client Protocols**, verify that all data protocols used by your remote access VPN clients are present. Add data protocols if necessary and then click **Next**.
5. In **Internet Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.
6. In **IP Address Assignment**, click **Automatic** if the VPN server should use DHCP to obtain IP addresses for remote access VPN clients. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the VPN clients to be reachable. When IP address assignment is complete, click **Next**.
7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and

authorization, click **Yes, I want to use a RADIUS server**, and then click **Next**.
   - In **RADIUS Server Selection**, configure the primary (mandatory) and secondary (optional) RADIUS servers and the shared secret, and then click **Next**.
8. Click **Finish**.
9. Start the Routing and Remote Access service when prompted.

By default, only 128 PPTP ports are configured on the WAN Miniport (PPTP) device. If you need more PPTP ports, configure the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in.

By default, only the MS-CHAP and MS-CHAPv2 protocols are enabled. If you are using smart cards or user certificates for authentication, select **Extensible Authentication Protocol (EAP)** check box from the **Authentication Methods** dialog box available from the **Security** tab on the properties of the VPN server in the Routing and Remote Access snap-in.

Intranet Network Infrastructure

Deploying the intranet network infrastructure for remote access VPN connections consists of the following:

- Configure routing on the VPN server.
- Verify name resolution and intranet reachability from the VPN server.
- Configure routing for off-subnet address pools.

Configuring routing on the VPN server

In order for your VPN servers to properly forward traffic to locations on the intranet, you must configure them with either static routes that summarize all the possible addresses used on the intranet or with routing protocols so that the VPN server can participate as a dynamic router and automatically add routes for intranet subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows 2000 Server Help. To configure the VPN server as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN server as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF".

Verifying name resolution and reachability from the VPN server

From each VPN server, verify that the VPN server can resolve names and successfully communicate with intranet resources by using the Ping command, Internet Explorer, and making drive and printer connections to known intranet servers.

Configuring routing for off-subnet address pools

If you configured any of the VPN servers with manual address pools and any of the pools are an off-subnet pool, you must ensure that the route(s) representing the off-subnet address pool(s) are present in your intranet routing infrastructure. You can ensure this by either adding static route(s) representing the off-subnet address pool(s) as static routes to the neighboring router(s) of the VPN server(s) and then using the routing protocol of your intranet to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the intranet interface of the VPN server.

Alternately, if you are using RIP or OSPF, you can configure the VPN servers using off-subnet address pools as

RIP or OSPF routers. For OSPF, you must configure the VPN server as an autonomous system boundary router (ASBR). For more information, see the topic titled "OSPF design considerations" in Windows 2000 Help.

Deploying VPN Clients

Deploying VPN clients for remote access VPN connections consists of the following:

- Manually configure VPN clients.
- Configure CM packages with CMAK.

Manually configuring VPN clients

If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows 2000 VPN clients, use the Make New Connection Wizard to create the Internet and VPN connections and link them together so that when you connect using the VPN connection, the Internet connection is already made. For Windows XP VPN clients, use the New Connection Wizard to create the Internet and VPN connections.

Configuring CM packages with CMAK

For a large number of VPN clients running different versions of Windows, you should use the CMAK to create and distribute customized Connection Manager packages for your users. For more information, see the topic titled "Before you start: Understanding Connection Manager and the Administration Kit" in Windows 2000 Server Help.

Deploying L2TP-based Remote Access

Deploying L2TP-based remote access VPN connections using Windows 2000 consists of the following:

- Deploy certificate infrastructure
- Deploy Internet infrastructure
- Deploy AAA infrastructure
- Deploy VPN servers
- Deploy intranet infrastructure
- Deploy VPN clients

Deploying Certificate Infrastructure

For L2TP-based VPN connections, a certificate infrastructure is required to issue computer certificates needed to negotiate authentication for IPSec. Additionally, a certificate infrastructure is also needed when you are using either smart cards or user certificates and EAP-TLS for user authentication. You must install a computer certificate on all VPN clients and VPN servers. If you are using EAP-TLS for user authentication, you must install a user certificate on all VPN clients and, if the authenticating server is a RADIUS server, a computer certificate on the RADIUS server.

For information about deploying a certificate infrastructure, see Appendix D, "Deploying a Certificate Infrastructure."

Deploying computer certificates

To install a computer certificate, a CA must be present to issue certificates. Once the CA is configured, you can install a computer certificate in the following ways:

1. By configuring the automatic allocation of computer certificates to computers in a Windows 2000 domain.
   This method allows a single point of configuration for the entire domain. All members of the domain automatically request the computer certificate through a group policy setting. To immediately obtain a computer certificate for a computer running Windows 2000 that is a member of the domain for which auto-enrollment is configured, restart the computer or type **secedit /refreshpolicy machine_policy** from a command prompt. To immediately obtain a computer certificate for a computer running Windows XP that is a member of the domain for which auto-enrollment is configured, restart the computer or type **gpupdate /target:computer** from a command prompt. To configure a Windows 2000 domain for automatic enrollment of computer certificates, see the topic titled "Configure automatic certificate allocation from an enterprise CA" in Windows 2000 Server Help.
2. By using the Certificates snap-in to request a computer certificate.
   If you are using a Windows 2000 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA using the Certificates snap-in. For more information, see the topics titled "Manage certificates for a computer" and "Request a certificate" in Windows 2000 Server Help.
3. By using the Certificates snap-in to import a computer certificate.
   If you have a certificate file that contains the computer certificate, you can import the computer certificate using the Certificates snap-in. For more information about importing a certificate using the Certificates snap-in, see the topic titled “Import a certificate” in Windows 2000 Server Help.
4. By executing a CAPICOM script that requests a computer certificate.
   In this method, each computer must execute a CAPICOM script that requests a computer certificate from the

issuing CA. CAPICOM is a COM client, supporting Automation, that performs cryptographic functions (the CryptoAPI) using Microsoft ActiveX® and COM objects. CAPICOM can be used via Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, see CAPICOM at http://msdn.microsoft.com/library/en-us/security/security/capicom_start_page.asp?frame=true.

Deploying smart cards

For information about deploying smart cards in Windows 2000, see the topic titled "Checklist: Deploying smart cards for logging on to Windows" in Windows 2000 Server Help.

Deploying user certificates

To install a user certificate, an issuing CA must be present to issue certificates. Once the issuing CA is configured, you can install a user certificate in the following ways:

1. By using a Web browser to request a user certificate.
   For more information about requesting a user certificate, see the topic titled “Submit a user certificate request via the Web” in Windows 2000 Server Help. The issuing CA must support Web enrollment of certificates. For example, if you are using a Windows 2000 enterprise CA as an issuing CA and the CA computer is also running Internet Information Services (IIS), you can use Web enrollment to request a user certificate.
2. By using the Certificates snap-in to request a user certificate.
   If you are using a Windows 2000 enterprise CA as an issuing CA, you can request a user certificate using the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled “Request a certificate” in Windows 2000 Server Help.
3. By importing a user certificate using the Certificates snap-in.
   If you have a certificate file that contains a user certificate, import the user certificate using the Certificates snap-in. For more information about requesting a user certificate using the Certificates snap-in, see the topic titled “Import a certificate” in Windows 2000 Server Help.
4. By executing a CAPICOM script that requests a user certificate.
   In this method, each user must execute a CAPICOM script that requests a user certificate from the issuing CA.

Deploying Internet Infrastructure

Deploying the Internet infrastructure for remote access VPN connections consists of the following:

- Place VPN servers in perimeter network or on the Internet.
- Install Windows 2000 Server on VPN servers and configure Internet interfaces.
- Add address records to Internet DNS.

Placing VPN servers in perimeter network or on the Internet

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between your intranet and the Internet. If so, configure packet filters on the firewall to allow L2TP/IPSec traffic to and from the IP address of the VPN servers' perimeter network interfaces. For more information, see Appendix A.

Installing Windows 2000 Server on VPN servers and configuring Internet interfaces

Install Windows 2000 Server on the VPN server computer and connect it to either the Internet or to perimeter network with one network adapter and connect it to the intranet with another network adapter. Without running the Routing and Remote Access Server Setup Wizard, the VPN server computer will not forward IP packets between the Internet and the intranet. For the connection connected to the Internet or the perimeter network, configure the TCP/IP protocol with a public IP address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the Internet.) Do not configure the connection with DNS server or WINS server IP addresses.

Adding address records to Internet DNS

To ensure that the name of the VPN server (for example, vpn.microsoft.com) can be resolved to its proper IP address, either add DNS address (A) records to your DNS server (if you are providing DNS name resolution for Internet users) or have your ISP add DNS address (Z) records to their DNS server(s) (if your ISP is providing DNS name resolution for Internet users). Verify that the name of the VPN server can be resolved to its public Internet IP address when connected to the Internet.

Deploying AAA Infrastructure

Deploying the AAA infrastructure for remote access VPN connections consists of the following:

- Configure Active Directory for user accounts and groups.
- Configure the primary IAS server on a domain controller.
- Configure the secondary IAS server on a different domain controller.

Configuring Active Directory for user accounts and groups

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all users that are making remote access connections have a corresponding user account. This includes employees, contractors, vendors, and business partners.
2. Set the remote access permission on user accounts to **Allow access** or **Deny access** to manage remote access by user. Or, to manage remote access by group, set the remote access permission on user accounts to **Control access through Remote Access Policy**.
3. Organize remote access users into the appropriate universal and nested groups in order to take advantage of group-based remote access policies. For more information, see the topic titled "Universal, global, and domain local groups" in Windows 2000 Server Help.

Configuring the primary IAS server on a domain controller

To configure the primary IAS server on a domain controller, do the following:

1. On the domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows 2000 Server Help.
2. Configure the IAS server computer (the domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help.
3. If the IAS server authenticates connection attempts for user accounts in other domains, verify that these domains have a two-way trust with the domain in which the IAS server computer is a member. Next,

configure the IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help. For more information about trust relationships, see the topic titled "Understanding domain trusts" in Windows 2000 Server Help.
If the IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.

4. Enable file logging for accounting and authentication events. For more information, see the topic titled "Configure log file properties" in Windows 2000 Server Help.
5. Add the VPN server(s) as RADIUS clients of the IAS server. For more information, see the topic titled "Add RADIUS clients" in Windows 2000 Server Help. For the IP address of each VPN server, use the intranet IP address assigned to the VPN server. If you are using names, use the internal name of the VPN server (this is not necessarily the same DNS name used by Internet clients). Use strong shared secrets.
6. Create remote access policies that reflect your remote access usage scenarios.
   For example, to configure a remote access policy that requires L2TP-based VPN connections for members of the Employees group to use EAP-TLS authentication and Triple-DES encryption, create a remote access policy with the following settings:

   Policy name: VPN connections

   Conditions:

   NAS-Port-Type matches **Virtual (VPN)**

   Tunnel-Type matches **Layer Two Tunneling Protocol**

   Windows-Groups matches **Employees** (example)

   Permission: **Grant remote access permission**

   Profile settings, **Authentication** tab:

   Select **Extensible Authentication Protocol** and the **Smart Card or other Certificate** EAP type. Clear all other check boxes.

   Profile settings, **Encryption** tab:

   Select the **Strongest** check box, and then clear all other check boxes.

7. If you have created new remote access policies, either delete the default remote access policy named **Allow access if dial-up permission is enabled**, or move it so that it is the last policy to be evaluated. For more information, see the topics titled "Delete a remote access policy" and "Change the policy evaluation order" in Windows 2000 Server Help.

Configuring the secondary IAS server on a different domain controller

To configure the secondary IAS server on a different domain controller, do the following:

1. On the other domain controller, install IAS as an optional networking component. For more information, see the topic titled "Install IAS" in Windows 2000 Server Help.
2. Configure the secondary IAS server computer (the other domain controller) to read the properties of user accounts in the domain. For more information, see the topic titled "Enable the IAS server to read user

objects in Active Directory" in Windows 2000 Server Help.

3. If the secondary IAS server authenticates connection attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the secondary IAS server computer is a member. Next, configure the secondary IAS server computer to read the properties of user accounts in other domains. For more information, see the topic titled "Enable the IAS server to read user objects in Active Directory" in Windows 2000 Server Help. For more information about trust relationships, see the topic titled "Understanding domain trusts" in Windows 2000 Server Help.
   If the secondary IAS server authenticates connection attempts for user accounts in other domains, and those domains do not have a two-way trust with the domain in which the secondary IAS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains.
4. Copy the configuration of the primary IAS server to the secondary IAS server. For more information, see the topic titled "Copy the IAS configuration to another server" in Windows 2000 Server Help.

Deploying VPN Servers

Deploying the VPN servers for remote access VPN connections consists of the following:

- Configure the VPN server's connection to the intranet.
- Run the Routing and Remote Access Server Setup Wizard.

Configuring the VPN server's connection to the intranet

For each VPN server, configure the connection connected to the intranet with a manual TCP/IP configuration consisting of IP address, subnet mask, intranet DNS servers, and intranet WINS servers. Note that you must not configure the default gateway on the intranet connection to prevent default route conflicts with the default route pointing to the Internet.

Running the Routing and Remote Access Server Setup Wizard

Run the Routing and Remote Access Server Setup Wizard to configure each Windows 2000 VPN server using the following steps:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Routing and Remote Access**.
2. Right-click your server name, and then click **Configure and Enable Routing and Remote Access**.
3. In **Common Configurations**, click **Virtual Private Network (VPN) server** and then click **Next**. If you want to use the VPN server computer as a network address translator (NAT), Web server, or other function, see Appendix B.
4. In **Remote Client Protocols**, verify that all data protocols used by your remote access VPN clients are present. Add data protocols if necessary and then click **Next**.
5. In **Internet Connection**, click the connection that corresponds to the interface connected to the Internet or your perimeter network, and then click **Next**.
6. In **IP Address Assignment**, click **Automatic** if the VPN server should use DHCP to obtain IP addresses for remote access VPN clients. Or, click **From a specified range of addresses** to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure in order for the VPN clients to be reachable. When IP address assignment is complete, click **Next**.
7. In **Managing Multiple Remote Access Servers**, if you are using RADIUS for authentication and

authorization, click **Yes, I want to use a RADIUS server**, and then click **Next**.
   - In **RADIUS Server Selection**, configure the primary (mandatory) and secondary (optional) RADIUS servers and the shared secret, and then click **Next**.
8. Click **Finish**.
9. Start the Routing and Remote Access service when prompted.

By default, only 128 L2TP ports are configured on the WAN Miniport (L2TP) device. If you need more L2TP ports, configure the **WAN Miniport (L2TP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in.

By default, only the MS-CHAP and MS-CHAPv2 protocols are enabled. If you are using smart cards or user certificates for authentication, select **Extensible Authentication Protocol (EAP)** check box from the **Authentication Methods** dialog box available from the **Security** tab on the properties of the VPN server in the Routing and Remote Access snap-in.

Intranet Network Infrastructure

Deploying the intranet network infrastructure for remote access VPN connections consists of the following:

- Configure routing on the VPN server.
- Verify name resolution and intranet reachability from the VPN server.
- Configure routing for off-subnet address pools.

Configuring routing on the VPN server

In order for your VPN servers to properly forward traffic to locations on the intranet, you must configure them with either static routes that summarize all the possible addresses used on the intranet or with routing protocols so that the VPN server can participate as a dynamic router and automatically add routes for intranet subnets to its routing table.

To add static routes, see the topic titled "Add a static route" in Windows 2000 Server Help. To configure the VPN server as a RIP router, see the topic titled "Configure RIP for IP". To configure the VPN server as an OSPF router, see the topics titled "OSPF design considerations" and "Configure OSPF".

Verifying name resolution and reachability from the VPN server

From each VPN server, verify that the VPN server can resolve names and successfully communicate with intranet resources by using the Ping command, Internet Explorer, and making drive and printer connections to known intranet servers.

Configuring routing for off-subnet address pools

If you configured any of the VPN servers with manual address pools and any of the pools are an off-subnet pool, you must ensure that the route(s) representing the off-subnet address pool(s) are present in your intranet routing infrastructure. You can ensure this by either adding static route(s) representing the off-subnet address pool(s) as static routes to the neighboring router(s) of the VPN server(s) and then using the routing protocol of your intranet to propagate the route to other routers. When you add the static route(s), you must specify that the gateway or next hop address is the intranet interface of the VPN server.

Alternately, if you are using RIP or OSPF, you can configure the VPN servers using off-subnet address pools as

RIP or OSPF routers. For OSPF, you must configure the VPN server as an autonomous system boundary router (ASBR). For more information, see the topic titled "OSPF design considerations" in Windows 2000 Help.

Deploying VPN Clients

Deploying VPN clients for remote access VPN connections consists of the following:

- Manually configure VPN clients.
- Configure CM packages with CMAK.

Manually configuring VPN clients

If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows 2000 VPN clients, use the Make New Connection Wizard to create the Internet and VPN connections and link them together so that when you connect using the VPN connection, the Internet connection is already made. For Windows XP VPN clients, use the New Connection Wizard to create the Internet and VPN connections.

Configuring CM packages with CMAK

For a large number of VPN clients running different versions of Windows, you should use the CMAK to create and distribute customized Connection Manager packages for your users. For more information, see the topic titled "Before you start: Understanding Connection Manager and the Administration Kit" in Windows 2000 Server Help.

Appendix A: Configuring Firewalls with a Windows 2000 VPN Server

The following are common configurations of firewalls with a VPN server:

- The VPN server is attached to the Internet and the firewall is between the VPN server and the intranet.
- The firewall is attached to the Internet and the VPN server is between the firewall and the intranet.
- Two firewalls are used—one between the VPN server and the intranet and one between the VPN server and the Intranet.

VPN Server in Front of the Firewall

To secure the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you need to configure PPTP or L2TP/IPSec input and output filters on the interface that corresponds to the connection to the Internet. Because IP routing is enabled on the Internet interface, if PPTP or L2TP/IPSec filters are not configured on the Internet interface, then any traffic received on the Internet interface is routed, which may forward unwanted Internet traffic to your intranet.

When the VPN server is in front of the firewall attached to the Internet, you need to add packet filters to the Internet interface that allow only VPN traffic to and from the IP address of the VPN server's Internet interface.

For inbound traffic, when the tunneled data is decrypted by the VPN server, it is forwarded to the firewall. The firewall in this configuration is acting as a filter for intranet traffic and can prevent specific resources from being accessed, scan data for viruses, perform intrusion detection, and other functions.

Because the only Internet traffic allowed on the intranet must pass through the VPN server, this approach also prevents the sharing of File Transfer Protocol (FTP) or Web intranet resources with non-VPN Internet users.

Figure 3 shows the VPN server in front of the firewall.

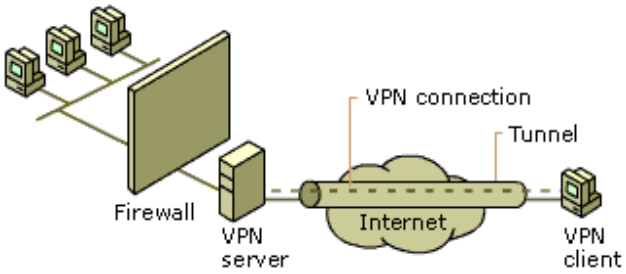

Figure 3 The VPN server in front of the firewall

The firewall is configured for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

For the Internet interface on the VPN server, configure the following input and output filters using the Routing and Remote Access snap-in. These filters are automatically configured by the Routing and Remote Access Server Setup Wizard in Windows 2000 Service Pack 2 and later.

Packet Filters for PPTP

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723.
  This filter allows PPTP tunnel maintenance traffic to the VPN server.
- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.
  This filter allows PPTP tunneled data to the VPN server.
- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] source port of 1723.
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. TCP [established] traffic is accepted only when the VPN server initiated the TCP connection.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723.
  This filter allows PPTP tunnel maintenance traffic from the VPN server.
- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47.
  This filter allows PPTP tunneled data from the VPN server.
- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] destination port of 1723.
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. TCP [established] traffic is sent only when the VPN server initiated the TCP connection.

Packet Filters for L2TP/IPSec

Configure the following input filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500.
  This filter allows Internet Key Exchange (IKE) traffic to the VPN server.
- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701.
  This filter allows L2TP traffic to the VPN server.

Configure the following output filters with the filter action set to **Drop all packets except those that meet the criteria below**:

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500.
  This filter allows IKE traffic from the VPN server.

- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701.
  This filter allows L2TP traffic from the VPN server.

There are no filters required for IPSec Encapsulating Security Protocol (ESP) traffic for the IP protocol of 50. The Routing and Remote Access service filters are applied after the IPSec components remove the ESP header.

VPN Server Behind the Firewall

In a more common configuration, the firewall is connected to the Internet and the VPN server is an intranet resource that is connected to the perimeter network, also known as a demilitarized zone (DMZ) or screened subnet. The perimeter network is an IP network segment that contains resources that are available to Internet users, such as Web and FTP servers. The VPN server has an interface on both the perimeter network and the intranet. In this approach, the firewall must be configured with input and output filters on its Internet interface that allow the passing of tunnel maintenance traffic and tunneled data to the VPN server. Additional filters can allow the passing of traffic to Web, FTP, and other types of servers on the perimeter network. For an added layer of security, the VPN server can also be configured with PPTP or L2TP/IPSec packet filters on its perimeter network interface.

The firewall in this configuration is acting as a filter for Internet traffic and can confine the incoming and outgoing traffic to the specific resources on the perimeter network, perform intrusion attempt detection, prevent denial of service attacks, and other functions.

Because the firewall does not have the encryption keys for each VPN connection, it can only filter on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server.

Figure 4 shows the VPN server behind the firewall on the perimeter network.

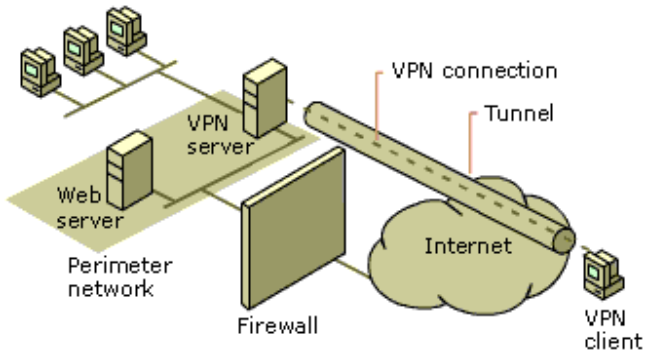

Figure 4 The VPN server behind the firewall on the perimeter network

For both the Internet and network perimeter interfaces on the firewall, configure the following input and output filters using the firewall's configuration software.

Packet Filters for PPTP

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

Filters on the Internet Interface

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).
  This filter allows PPTP tunnel maintenance traffic to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).
  This filter allows PPTP tunneled data to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

Configure the following output filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).
  This filter allows PPTP tunnel maintenance traffic from the VPN server.
- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).
  This filter allows PPTP tunneled data from the VPN server.
- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Filters on the Perimeter Network Interface

Configure the following input filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).
  This filter allows PPTP tunnel maintenance traffic from the VPN server.
- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).
  This filter allows PPTP tunneled data from the VPN server.
- Source IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all

traffic from the VPN server to TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB).
  This filter allows PPTP tunnel maintenance traffic to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F).
  This filter allows PPTP tunneled data to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB).
  This filter is required only when the VPN server is acting as a VPN client (a calling router) in a router-to-router VPN connection. This filter should only be used in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet using this port.

Packet Filters for L2TP/IPSec

Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

Filters on the Internet Interface

Configure the following input packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).
  This filter allows IKE traffic to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).
  This filter allows IPSec ESP traffic to the VPN server.

Configure the following output packet filters on the Internet interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).
  This filter allows IKE traffic from the VPN server.
- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).
  This filter allows IPSec ESP traffic from the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

Filters on the Perimeter Network Interface

Configure the following input packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Source IP address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4).
  This filter allows IKE traffic from the VPN server.
- Source IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).
  This filter allows IPSec ESP traffic from the VPN server.

Configure the following output packet filters on the perimeter network interface of the firewall to allow the following types of traffic:

- Destination IP address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4).
  This filter allows IKE traffic to the VPN server.
- Destination IP address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32).
  This filter allows IPSec ESP traffic to the VPN server.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPSec ESP payload.

VPN Server Between Two Firewalls

Another configuration is when the VPN server computer in placed on the perimeter network between two firewalls. The Internet firewall, the firewall between the Internet and the VPN server, filters all Internet traffic from all Internet clients. The intranet firewall, the firewall between the VPN server and the intranet, filters intranet traffic from VPN clients.

Figure 5 shows the VPN server between two firewalls on the perimeter network.

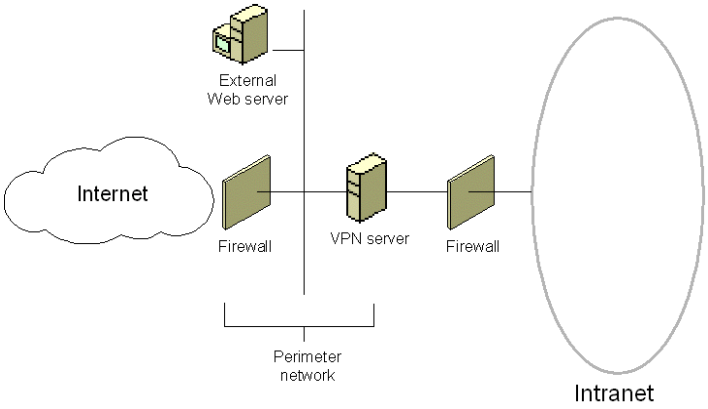

Figure 5 The VPN server between two firewalls on the perimeter network

In this configuration:

- Configure your Internet firewall and VPN server with the packet filters as described in the "VPN Server Behind the Firewall" section.
- Configure your intranet firewall for the appropriate rules for intranet traffic to and from VPN clients

according to your network security policies.

Appendix B: Alternate Configurations

This section provides information about common alternate configurations for a Windows 2000 VPN server. The most common configuration is described in the "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" sections of this paper and whose principal characteristics are the following:

- The VPN server has multiple network adapters—at least one connected to the intranet and at least one connected to the Internet.
- The VPN server has static public IP addresses assigned to its Internet interfaces.
- The VPN server is only acting as a security gateway providing remote access to the intranet. The VPN server is not hosting any other Internet services such as NAT or Web services.

The two other most common configurations are the following:

1. The VPN server computer is performing other functions such as network address translation or Web hosting.
2. The VPN server computer has a single network adapter and its public IP address is published by a firewall.

The following sections detail the changes to make in the deployment of a VPN server to accommodate these additional common configurations.

Multiple Internet Function VPN Server

In this configuration, the VPN server's principal characteristics are the following:

- The VPN server has multiple network adapters—at least one connected to the intranet and at least one connected to the Internet.
- The VPN server has static public IP addresses assigned to its Internet interfaces.
- The VPN server is acting as a security gateway providing remote access to the intranet and is hosting any other Internet services such as NAT or Web hosting.

In this configuration, you can follow the procedures as described in the "Deploying PPTP-based Remote Access" and "Deploying L2TP-based Remote Access" sections of this paper except that when you run the Routing and Remote Access Server Setup Wizard, you select from the list of **Common Configurations**, do not choose **Virtual Private Network (VPN) server**. Instead, select **Remote access server**. You are prompted to select an interface over which DHCP, DNS, and WINS configuration is obtained, to determine how you want to assign IP addresses to remote access clients, and to configure RADIUS.

When you select **Remote access server**, only five PPTP and L2TP ports are configured. For additional ports, configure the properties of the **WAN Miniport (PPTP)** and **WAN Miniport (L2TP)** devices from the properties of the **Ports** object in the Routing and Remote Access snap-in.

By selecting **Remote access server** in the wizard, PPTP and L2TP packet filters are not configured on the Internet interface of the VPN server computer. Whether you have to manually configure these filters depends on whether the VPN server computer is also hosting NAT.

- If NAT is needed on the VPN server computer, do not configure PPTP and L2TP packet filters or packet filters for other types of traffic. If you configure PPTP and L2TP packet filters on the Internet interface, NAT cannot function. Even though you do not configure any packet filters on the Internet interface of the VPN server computer, the function of the NAT discards any traffic from the Internet that does not

correspond to traffic requested by intranet clients.

- If NAT is not needed on the VPN server computer, you can configure PPTP and L2TP packet filters and other types of filters for additional services hosted by the VPN server computer. For example, if the VPN server computer is also hosting a Web site, then filters should be added to allow traffic to and from the public IP address of the VPN server computer and TCP port 80.

Single-Adapter VPN Server

In this configuration, the VPN server computer has only a single network adapter and VPN clients are accessing services hosted on the VPN server computer. If the VPN server computer has only a single network adapter and is configured with a public IP address, all traffic to and from the services running on the VPN server computer are sent as clear text outside the VPN tunnel. For more information about why this happens, see "Routing and multi-use VPN servers" in this paper.

The only way a single adapter VPN server can work properly is if it is behind a firewall that is providing a publishing and translation service for the VPN server. The firewall publishes or makes known on the Internet a static public IP address for the VPN server. When VPN packets are sent to this published IP address, the firewall translates the address of the packet to a private or other public address by which the VPN server is known on the intranet.

Figure 6 shows an example of the published and actual addresses of a VPN server in this configuration.

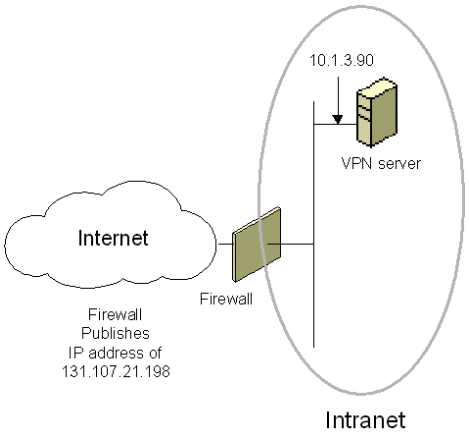

Figure 6 The single-adapter VPN server configuration

Just as in "Routing and multi-use VPN servers" in this paper, a VPN client uses the Internet DNS to resolve the VPN server's name to its published public IP address. After the VPN connection is made, the intranet DNS and WINS infrastructures resolve the VPN server's name to its actual intranet address. One limitation to this configuration is that only PPTP is supported. Because the firewall is translating addresses, IPSec-protected L2TP traffic cannot traverse the firewall.

The VPN server is configured according to "Deploying PPTP-based Remote Access" in this paper with its intranet interface acting as an Internet interface. The firewall is configured to:

- Publish the name and public IP address of the VPN server on the Internet.

- Translate PPTP traffic sent to the public IP address of the VPN server to the intranet interface of the VPN server computer.
- Discard all traffic except PPTP traffic going to and from the VPN server computer.

Appendix C: Setting up a VPN test lab

This section provides detailed information about how you can use five computers to create a test lab with which to configure and test the virtual private network (VPN) features in Windows 2000. These instructions are designed to take you through a set of tasks that expose you to VPN connections and their functionality. Beyond the set of tasks, you can use these instructions to create a functioning VPN configuration. You can then use this configuration to experiment with VPN features and functionality in prior to deployment on your production network.

This section covers:

- Setting up the infrastructure
- Virtual private network test lab tasks

Setting up the infrastructure

The infrastructure for the VPN test lab network consists of five computers performing the following services:

- A computer running Windows 2000 that is acting as a domain controller, a Domain Name System (DNS) server, and a certification authority (CA). This computer is named DC1.
- A computer running Windows 2000 that is acting as a Remote Authentication Dial-in User Service (RADIUS) server. This computer is named IAS1.
- A computer running Windows 2000 that is acting as a Web server and file sharing server. This computer is named IIS1.
- A computer running Windows 2000 that is acting as a VPN server. This computer is named VPN1. VPN1 has two network adapters installed.
- A computer running Windows 2000 that is acting as a VPN client. This computer is named CLIENT1.

Figure 7 shows the configuration of the VPN test lab.

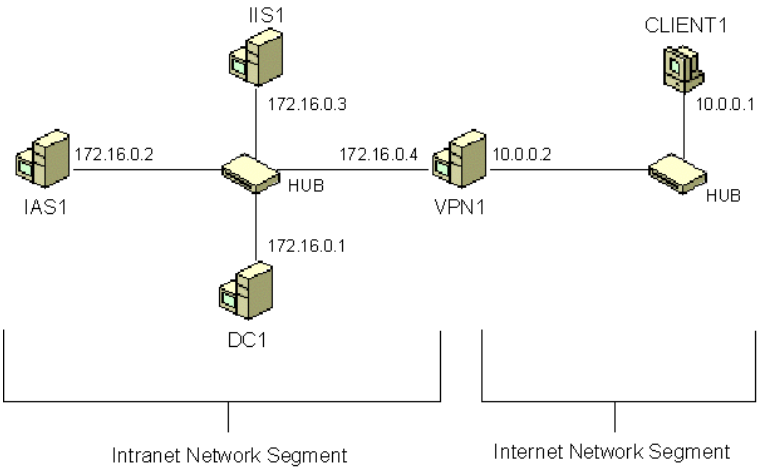

Figure 7. Configuration of the VPN test lab

There is a network segment representing a corporate intranet and a network segment representing the Internet. All

computers on the corporate intranet are connected to a common hub or Layer 2 switch. All computers on the Internet are connected to a separate common hub or Layer 2 switch. Private addresses are used throughout the test lab configuration. The private network of 172.16.0.0/24 is used for the intranet. The private network of 10.0.0.0/24 is used for the simulated Internet.

Each computer is manually configured with the appropriate IP address, subnet mask, and DNS server IP address. There are no Dynamic Host Configuration Protocol (DHCP) or Windows Internet Name Service (WINS) servers present.

The following sections describe the configuration for each of the computers in the test lab. To reconstruct this test lab, configure the computers in the order presented.

Note: The following instructions are for configuring a test lab using a minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor is it designed to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

DC1

DC1 is a computer running Windows 2000 that is providing the following services:

- A domain controller for the testlab.microsoft.com domain.
- A DNS server for the testlab.microsoft.com DNS domain.
- The enterprise root certification authority (CA) for the testlab.microsoft.com domain.

To configure DC1 for these services, perform the following steps.

1. Install Windows 2000 as a stand-alone server.
2. Configure the TCP/IP protocol with the IP address of 172.16.0.1 and the subnet mask of 255.255.255.0.
3. Run dcpromo for a new domain called testlab.microsoft.com in a new forest. Install the DNS service when prompted.
4. Install the Certificate Services component as an enterprise root certification authority.
5. Configure the testlab.microsoft.com domain for automatic enrollment of computer certificates.

IAS1

IAS1 is a computer running Windows 2000 that is providing RADIUS authentication, authorization, and accounting for VPN1 (the VPN server computer).

To configure IAS1 as a RADIUS server, perform the following steps:

1. On DC1, add a computer account for the IAS1 computer.
2. Install Windows 2000 as a stand-alone server.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.2, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
4. Join IAS1 to the testlab.microsoft.com domain.
5. Install the Internet Authentication Service.

IIS1

IIS1 is a computer running Windows 2000 and the Internet Information Service. It is providing Web server services for intranet clients. To configure IIS1 as a Web server, perform the following steps:

1. On DC1, add a computer account for the IIS1 computer.
2. Install Windows 2000 as a stand-alone server.
3. Configure the TCP/IP protocol with the IP address of 172.16.0.3, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
4. Join IIS1 to the testlab.microsoft.com domain.
5. Install the Internet Information Service.
6. To determine whether the Web server is working correctly, run Internet Explorer on IAS1. When prompted by the Internet Connection wizard, configure the wizard for a LAN connection. In Internet Explorer, in **Address**, type http://IIS1.testlab.microsoft.com/win2000.gif. You should see a Windows 2000 graphic.
7. In Windows Explorer, share the root directory of Local Disk (C:)using the share name ROOT to the group **Everyone** with full access.

VPN1

VPN1 is a computer running Windows 2000 that is providing VPN server services for Internet-based VPN clients. To configure VPN1 as a VPN server, perform the following steps:

1. On DC1, add a computer account for VPN1.
2. Install Windows 2000 as a stand-alone server.
3. For the intranet local area connection, configure the TCP/IP protocol with the IP address of 172.16.0.4, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
4. For the Internet local area connection, configure the TCP/IP protocol with the IP address of 10.0.0.2 and the subnet mask of 255.255.255.0.
5. Join VPN1 to the testlab.microsoft.com domain.
6. Configure and enable the Routing and Remote Access service. In the Routing and Remote Access Server Setup Wizard, select **Virtual private network (VPN) server** from the list of common configurations. When prompted for IP address assignment, select **From a specified range of addresses** and configure the range 172.16.0.248 to 172.16.0.255. Do not configure RADIUS authentication.

CLIENT1

CLIENT1 is a computer running Windows 2000 that is acting as a VPN client and gaining remote access to intranet resources across the simulated Internet. To configure CLIENT1 as a VPN client, perform the following steps:

1. On DC1, add a computer account for CLIENT1.
2. Connect CLIENT1 to the intranet network segment.
3. On CLIENT1, install Windows 2000 as a workgroup computer.
4. Configure the TCP/IP protocol with the IP address of 172.16.0.5, the subnet mask of 255.255.255.0, and the DNS server IP address of 172.16.0.1.
5. Join CLIENT1 to the testlab.microsoft.com domain.
6. Configure the TCP/IP protocol with the IP address of 10.0.0.1, the subnet mask of 255.255.255.0, and no DNS server IP address.
7. Shut down the CLIENT1 computer.
8. Disconnect the CLIENT1 computer from the intranet network segment and connect it to the simulated

Internet network segment.

9. Restart the CLIENT computer and log on using the cached credentials of the testlab.microsoft.com administrator account.

VPN test lab tasks

The following tasks are designed to take you through the most common elements of remote access VPN support with Windows 2000:

- PPTP-based remote access
- L2TP-based remote access
- RADIUS authentication and accounting
- Remote access policies for different types of VPN connections

PPTP-based remote access

To create a PPTP-based remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following steps:

Create a user account

On DC1, use the Active Directory Users and Computers snap-in to create a user account named PPTPUser with a password. Set the remote access permission on the **Dial-in** tab to **Allow access**.

Create the PPTP connection

1. On CLIENT1, use the Make New Connection Wizard to create a new VPN connection named PPTPtoCorpnet, using the VPN server IP address of 10.0.0.2.
2. Right-click the new **PPTPtoCorpnet** connection, and then click **Properties**.
3. Click the **Networking** tab, and then in **Type of VPN**, click **Point to Point Tunneling Protocol**.
4. Click **OK** to save changes to the **PPTPtoCorpnet** connection.

Make the PPTP connection

1. On CLIENT1, double-click the **PPTPtoCorpnet** connection.
2. In the PPTPtoCorpnet dialog box, type **PPTPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.
3. Click **Connect**.

Access Web server and file share on the intranet

1. On CLIENT1, run Internet Explorer.
2. When prompted by the Internet Connection Wizard, configure it for a LAN connection.
3. In Internet Explorer, in **Address**, type **http://IIS1.testlab.microsoft.com/win2000.gif**. You should see a Windows 2000 graphic.
4. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the Local Drive (C:) on IIS1.

Disconnect the PPTP connection

On CLIENT1, right-click the PPTPtoCorpnet connection and then click **Disconnect**.

L2TP-based remote access

To create an L2TP-based remote access VPN connection between CLIENT1 and VPN1 and test whether intranet resources are available, perform the following:

Create a user account

On DC1, use Active Directory Users and Computers to create a user account named L2TPUser with a password. Set the remote access permission on the **Dial-in** tab to **Allow access**.

Create the L2TP connection

1. On CLIENT1, use the Make New Connection Wizard to create a new VPN connection named L2TPtoCorpnet, using the VPN server IP address of 10.0.0.2.
2. Right-click the new **L2TPtoCorpnet** connection, and then click **Properties**.
3. Click the **Networking** tab, and then in **Type of VPN**, click **Layer 2 Tunneling Protocol**.
4. Click **OK** to save changes to the **L2TPtoCorpnet** connection.

Make the L2TP connection

1. On CLIENT1, double-click the **L2TPtoCorpnet** connection.
2. In the L2TPtoCorpnet dialog box, type **L2TPUser@testlab.microsoft.com** as the user name, type the password, and then select the **Save this user name and password to use when** check box.
3. Click **Connect**.

Access Web server and file share on the intranet

1. On CLIENT1, run Internet Explorer.
2. In Internet Explorer, in **Address**, type **http://IIS1.testlab.microsoft.com/win2000.gif**. You should see a Windows 2000 graphic.
3. On CLIENT1, click **Start**, click **Run**, type **\\IIS1\ROOT**, and then click **OK**. You should see the contents of the Local Drive (C:) on IIS1.

Disconnect the L2TP connection

On CLIENT1, right-click the L2TPtoCorpnet connection and then click **Disconnect**.

RADIUS authentication and accounting

To configure RADIUS authentication and accounting for VPN connections, perform the following:

Configure IAS1 for VPN1 as a RADIUS client

On IAS1, add VPN1 as a RADIUS client using the IP address of 172.16.0.4 and a shared secret. To add a RADIUS client, right-click the **Clients** folder and click **New Client** in the Internet Authentication Service snap-in.

Configure IAS1 to log authentication events

On IAS1, enable the logging of accounting and authentication requests from **Settings** tab in the properties of the **Local File** object in the **Remote Access Logging** folder in the Internet Authentication Service snap-in.

Configure VPN1 for IAS1 as a RADIUS server

On VPN1, add IAS1 as a RADIUS server for both the authentication and accounting provider at the IP address of 172.16.0.2 and the shared secret. To configure the Routing and Remote Access service for RADIUS, obtain properties on the VPN server and click the **Security** tab. For RADIUS authentication, select **RADIUS authentication** as the authentication provider and click **Configure** to add IAS1 as the RADIUS server. For RADIUS accounting, select **RADIUS accounting** as the accounting provider and click **Configure** to add IAS1 as the RADIUS server.

Make PPTP and L2TP connections

1. On CLIENT1, make a PPTP connection with VPN1 using the PPTPtoCorpnet connection.
2. Disconnect the PPTP connection.
3. On CLIENT1, make an L2TP connection with VPN1 using the L2TPtoCorpnet connection.
4. Disconnect the L2TP connection.

Check the system event log for RADIUS events

On IAS1, use Event Viewer to view IAS events in the system event log for the PPTP and L2TP connections that were recently created using CLIENT1.

Check RADIUS authentication and accounting logs

On IAS1, use Windows Explorer to open the *SystemRoot*\System32\Logfiles\Iaslog.log file. Note the authentication and accounting entries for the PPTP and L2TP connections that were recently created using CLIENT1.

Remote access policies for different types of VPN connections

To create remote access policies for different types of VPN connections, do the following:

Create separate remote access policies for PPTP and L2TP connections

1. On IAS1, create a new remote access policy with the following settings:
   Policy name: PPTP connections

   Conditions:

   NAS-Port-Type matches **Virtual (VPN)**

   Tunnel-Type matches **Point-to-Point Tunneling Protocol (PPTP)**

   Permission: Grant remote access permission

   Profile settings, **IP** tab:

   From client packet filter:

   Filter action: **Deny all traffic except those listed below**

Destination network, IP address: 172.16.0.1

Destination network, Subnet mask: 255.255.255.255

Protocol: Any

To client packet filter:

Filter action: **Deny all traffic except those listed below**

Source network, IP address: 172.16.0.1

Destination network, Subnet mask: 255.255.255.255

Protocol: Any

2. Create a new custom remote access policy with the following settings:
   Policy name: L2TP connections

   Conditions:

   NAS-Port-Type matches **Virtual (VPN)**

   Tunnel-Type matches **Layer Two Tunneling Protocol (L2TP)**

   Permission: Grant remote access permission

   Profile settings, **IP** tab:

   From client packet filter:

   Filter action: Deny all traffic except those listed below

   Destination network, IP address: 172.16.0.2

   Destination network, Subnet mask: 255.255.255.255

   Protocol: Any

   To client packet filter:

   Filter action: Deny all traffic except those listed below

   Source network, IP address: 172.16.0.2

   Destination network, Subnet mask: 255.255.255.255

   Protocol: Any

Make a PPTP connection and test connectivity

1. On CLIENT1, make a VPN connection with VPN1 using the PPTPtoCorpnet connection.
2. Use the ping command to ping DC1 at its IP address of 172.16.0.1.
3. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.
   This command fails because packet filtering for all connections that match the **PPTP connections** policy allows only traffic sent to and from the IP address of 172.16.0.1.
4. Disconnect the PPTPtoCorpnet connection.

Make an L2TP connection and test connectivity

1. On CLIENT1, make a VPN connection with VPN1 using the L2TPtoCorpnet connection.
2. Use the ping command to ping IAS1 at its IP address of 172.16.0.2.
3. Use the ping command to ping DC1 at its IP address of 172.16.0.1.
   This command fails because packet filtering for all connections that match the **L2TP connections** policy allows only traffic sent to and from the IP address of 172.16.0.2.
4. Disconnect the L2TPtoCorpnet connection.

Check the system event log for IAS events

On IAS1, use Event Viewer to view the IAS events in the system event log for the PPTP and L2TP connections that were recently created by CLIENT1. Note that the authentication event message text contains the name of the remote access policy that accepted the connection.

Appendix D: Troubleshooting

Troubleshooting tools

Windows 2000 provides the following tools to troubleshoot VPN connections:

- TCP/IP Troubleshooting Tools
- Authentication and Accounting Logging
- Event Logging
- IAS Event Logging
- PPP Logging
- Tracing
- Network Monitor

TCP/IP Troubleshooting Tools

The Ping, Tracert, and Pathping tools use ICMP Echo and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The **route print** command can be used to display the IP routing table. Alternately, on the VPN server, you can use the **netsh routing ip show rtmroutes** command or the Routing and Remote Access snap-in. The Nslookup tool can be used to troubleshoot DNS and name resolution issues.

In addition to the normal TCP/IP tools, use the Netdiag tool to test and display your network configuration.

Authentication and Accounting Logging

A VPN server running Windows 2000 supports the logging of authentication and accounting information for remote access VPN connections in local logging files when Windows authentication or Windows accounting is enabled. This logging is separate from the events recorded in the system event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting remote access policy issues. For each authentication attempt, the name of the remote access policy that either accepted or rejected the connection attempt is recorded.

Enable authentication and accounting logging from the **Settings** tab on the properties of the **Local File** object in the **Remote Access Logging** folder in the Routing and Remote Access snap-in (if the Routing and Remote Access service is configured for Windows authentication and accounting) or the Internet Authentication Service snap-in (if the Routing and Remote Access service is configured for RADIUS authentication and accounting and the RADIUS server is an IAS server)

The authentication and accounting information is stored in a configurable log file or files stored in the *SystemRoot*\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis.

If the VPN server is configured for RADIUS authentication and accounting and the RADIUS server is a computer running Windows 2000 and IAS, the authentication and accounting logs are stored in the *SystemRoot*\System32\LogFiles folder on the IAS server computer.

Event Logging

On the **Event Logging** tab in the properties of a VPN server in the Routing and Remote Access snap-in, there are

four levels of logging. Select **Log the maximum amount of information**, and then try the connection again. After the connection fails, check the system event log for events logged during the connection process. After you are done viewing remote access events, select the **Log errors and warnings option** on the **Event logging** tab to conserve system resources.

IAS Event Logging

If your VPN servers are configured for RADIUS authentication and your RADIUS servers are computers running Windows 2000 Server and IAS, check the system event log for IAS events for rejected or accepted connection attempts. IAS system event log entries contain a lot of information on the connection attempt including the name of the remote access policy that accepted or rejected the connection attempt. IAS event logging for rejected or accepted connection attempts is enabled by default and configured from the **Service** tab from the properties of an IAS server in the Internet Authentication Service snap-in.

PPP logging

PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, select the **Enable Point-to-Point Protocol (PPP) logging** option on the **PPP** tab on the properties of a remote access server.

The PPP log in Windows NT 4.0 has been replaced by the tracing function. To duplicate the PPP log, you need to enable file tracing for the PPP key. By default, the PPP log is stored as the Ppp.log file in the *SystemRoot*\Tracing folder.

Tracing

The Windows 2000 Routing and Remote Access service has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable the components in Windows 2000 Server to log tracing information to files using the Netsh command or through the registry.

Enabling Tracing with Netsh

You can use the Netsh command to enable and disable tracing for specific components or for all components. To enable and disable tracing for a specific component, use the following syntax:

netsh ras set tracing *Component* **enabled|disabled**

where *Component* is a component in the list of Routing and Remote Access service components found in the Windows 2000 registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, the command is:

netsh ras set tracing rasauth enabled

To enable tracing for all components, use the following command:

netsh ras set tracing * enabled

Enabling Tracing through the Registry

The tracing function can also be configured by changing settings in the Windows 2000 registry under:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing

You can enable tracing for each Routing and Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing and Remote Access service is running. Each component is capable of tracing and appears as a subkey under the preceding registry key.

To enable tracing for each component, you can configure the following registry value entries for each protocol key:

EnableFileTracing REG_DWORD *Flag*

You can enable logging tracing information to a file by setting **EnableFileTracing** to 1. The default value is 0.

FileDirectory REG_EXPAND_SZ *Path*

You can change the default location of the tracing files by setting **FileDirectory** to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the *SystemRoot*\Tracing folder.

FileTracingMask REG_DWORD *LevelOfTracingInformationLogged*

FileTracingMask determines how much tracing information is logged to the file. The default value is 0xFFFF0000.

MaxFileSize REG_DWORD *SizeOfLogFile*

You can change the size of the log file by setting different values for **MaxFileSize**. The default value is 0x10000 (64K).

Notes: Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

Tracing information can be complex and very detailed. Most of the time this information is useful only to Microsoft support professionals or to network administrators who are very experienced with the Routing and Remote Access service. Tracing information can be saved as files and sent to Microsoft support for analysis.

Network Monitor

Use Network Monitor, a packet capture and analysis tool supplied with Windows 2000 Server, to capture and view the traffic sent between a VPN server and VPN client during the VPN connection process and during data transfer. You cannot interpret the encrypted portions of VPN traffic with Network Monitor. Network Monitor is installed as an optional networking component.

The proper interpretation of the remote access and VPN traffic with Network Monitor requires an in-depth understanding of PPP, PPTP, IPSec, and other protocols. Network Monitor captures can be saved as files and sent to Microsoft support for analysis.

Troubleshooting remote access VPNs

Remote access VPN problems typically fall into the following categories:

- Connection attempt is rejected when it should be accepted.
- Connection attempt is accepted when it should be rejected.

- Unable to reach locations beyond the VPN server.
- Unable to establish a tunnel.

Use the following troubleshooting tips to isolate the configuration or infrastructure problem causing the stated VPN problem.

Connection attempt is rejected when it should be accepted

- Using the Ping command, verify that the host name is being resolved to its correct IP address. The ping itself might not be successful due to packet filtering that is preventing the delivery of ICMP messages to and from the VPN server.
- Verify that the VPN client's credentials, consisting of user name, password, and domain name, are correct and can be validated by the VPN server.
- Verify that the user account of the VPN client is not locked out, expired, disabled, or that the time the connection is being made does not correspond to the configured logon hours. If the password on the account has expired, verify that the remote access VPN client is using MS-CHAP v1 or MS-CHAP v2. MS-CHAP v1 and MS-CHAP v2 are the only authentication protocols provided with Windows 2000 that allow you to change an expired password during the connection process.
  For an administrator-level account whose password has expired, reset the password using another administrator-level account.
- Verify that the user account has not been locked out due to remote access account lockout.
- Verify that the Routing and Remote Access service is running on the VPN server.
- Verify that the VPN server is enabled for remote access from the **General** tab on the properties of a VPN server in the Routing and Remote Access snap-in.
- Verify that the **WAN Miniport (PPTP)** and **WAN Miniport (L2TP)** devices are enabled for inbound remote access from the properties of the **Ports** object in the Routing and Remote Access snap-in.
- Verify that the VPN client, the VPN server, and the remote access policy corresponding to VPN connections are configured to use at least one common authentication method.
- Verify that the VPN client and the remote access policy corresponding to VPN connections are configured to use at least one common encryption strength.
- Verify that the parameters of the connection have permission through remote access policies.
  In order for the connection to be established, the parameters of the connection attempt must:
  - Match all of the conditions of at least one remote access policy.
  - Be granted remote access permission through the user account (set to **Allow access**), or if the user account has the **Control access through Remote Access Policy** option selected, the remote access permission of the matching remote access policy must have the **Grant remote access permission** option selected.
  - Match all the settings of the profile.
  - Match all the settings of the dial-in properties of the user account.

  To obtain the name of the remote access policy that rejected the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name.
- If you are logged on using an account with domain administrator permissions when you run the Routing and Remote Access Server Setup Wizard, it automatically adds the computer account of the **RAS and IAS Servers** domain-local security group. This group membership allows the VPN server computer to access user account information. If the VPN server is unable to access user account information, verify

that:

  - The computer account of the VPN server computer is a member of the **RAS and IAS Servers** security group for all the domains that contain user accounts for which the VPN server is authenticating remote access. You can use the **netsh ras show registeredserver** command at the command prompt to view the current registration. You can use the **netsh ras add registeredserver** command to register the server in a domain in which the VPN server is a member or other domains. Alternately, you or your domain administrator can add the computer account of the VPN server computer to the **RAS and IAS Servers** security group of all the domains that contain user accounts for which the VPN server is authenticating remote access.
  - If you add or remove the VPN server computer to the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Windows 2000 caches Active Directory information). For the change to take effect immediately, you need to restart the VPN server computer.
- For a VPN server that is a member server in a mixed-mode or native-mode Windows 2000 domain that is configured for Windows authentication, verify that:
  - The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.
  - The **RAS and IAS Servers** security group has **Read** permission to the **RAS and IAS Servers Access Check** object.
- Verify that the LAN protocols (TCP/IP, IPX, NetBEUI) used by the VPN client are enabled for remote access on the VPN server.
- Verify that all of the PPTP or L2TP ports on the VPN server are not already being used. If necessary, change the number of PPTP to L2TP ports from the properties of the **Ports** object in the Routing and Remote Access snap-in to allow more concurrent connections.
- Verify that the tunneling protocol of the VPN client is supported by the VPN server.
  By default, Windows 2000 remote access VPN clients have the **Automatic** server type option selected, which means that they try to establish a L2TP/IPSec-based VPN connection first, then they try a PPTP-based VPN connection. If either the **Point to Point Tunneling Protocol (PPTP)** or **Layer-2 Tunneling Protocol (L2TP)** server type option is selected, verify that the selected tunneling protocol is supported by the VPN server.

  By default, Windows XP remote access VPN clients have the **Automatic** VPN type option selected, which means that they try to establish a PPTP -based VPN connection first, then they try a L2TP/IPSec-based VPN connection. If either the **PPTP VPN** or **L2TP IPSec VPN** type is selected, verify that the selected tunneling protocol is supported by the VPN server.

  Depending on your selections when running the Routing and Remote Access Server Setup Wizard, a Windows 2000 Server–based computer running the Routing and Remote Access service is a PPTP and L2TP server with five or 128 L2TP ports and five or 128 PPTP ports. To create a PPTP-only server, set the number of L2TP ports to zero. To create an L2TP-only server, set the number of PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the **WAN Miniport (PPTP)** device from the properties of the **Ports** object in the Routing and Remote Access snap-in.

- For L2TP/IPSec connections, verify that computer certificates, also known as machine certificates, are installed on the VPN client and the VPN server.
- If the VPN server is configured with static IP address pools, verify that there are enough addresses. If all of the addresses in the static pools have been allocated to connected VPN clients, the VPN server is unable to allocate an IP address for TCP/IP-based connections, and the connection attempt is rejected.
- If the VPN client is configured to request its own IPX node number, verify that the VPN server is

configured to allow IPX clients to request their own IPX node number.

- If the VPN server is configured with a range of IPX network numbers, verify that the IPX network numbers in the range are not being used elsewhere on your IPX internetwork.
- Verify the configuration of the authentication provider. The VPN server can be configured to use either Windows or RADIUS to authenticate the credentials of the VPN client.
  - For RADIUS authentication, verify that the VPN server computer can communicate with the RADIUS server.
  - For a VPN server that is a member of a Windows 2000 native-mode domain, verify that the VPN server has joined the domain.
  - For a Windows NT version 4.0 Service Pack 4 and later VPN server that is a member of a Windows 2000 mixed mode domain or a Windows 2000 VPN server that is a member of a Windows NT 4.0 domain that is accessing user account properties for a user account in a trusted Windows 2000 domain, verify that the **Everyone** group is added to the **Pre-Windows 2000 Compatible Access** group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, issue the net **localgroup "Pre-Windows 2000 Compatible Access" everyone /add** command on a domain controller computer and then restart the domain controller.
  - For a Windows NT version 4.0 Service Pack 3 and earlier VPN server that is a member of a Windows 2000 mixed-mode domain, verify that the **Everyone** group has been granted list contents, read all properties, and read permissions to the root node of your domain and all sub-objects of the root domain.
- For PPTP connections using MS-CHAP v1 and attempting to negotiate 40-bit MPPE encryption, verify that the user's password is not larger than 14 characters.

Connection attempt is accepted when it should be rejected

- Verify that the remote access permission on the user account is set to either **Deny access** or **Control access through Remote Access Policy**. If set to the latter, verify that the first matching remote access policy's remote access permission is set to **Deny remote access permission**. To obtain the name of the remote access policy that accepted the connection attempt, scan the accounting log for the entry corresponding to the connection attempt for the policy name.
- If you have created a remote access policy to explicitly reject all connections, verify the policy conditions, remote access permission, and profile settings.

Unable to reach locations beyond the VPN server

- Verify that either the protocol is enabled for routing or that dial-in clients are allowed to access the entire network for LAN protocols being used by the VPN clients.
- Verify the IP address pools of the VPN server.
  If the VPN server is configured to use a static IP address pool, verify that the routes to the range of addresses defined by the static IP address pools are reachable by the hosts and routers of the intranet. If not, then IP route consisting of the VPN server static IP address pools, as defined by the IP address and mask of the range, must be added to the routers of the intranet or enable the routing protocol of your routed infrastructure on the VPN server. If the routes to the remote access VPN client subnets are not present, remote access VPN clients cannot receive traffic from locations on the intranet. Routes for the subnets are implemented either through static routing entries or through a routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP).

If the VPN server is configured to use DHCP for IP address allocation and no DHCP server is available, the VPN server assigns addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the VPN server is attached is also using APIPA addresses.

If the VPN server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the VPN server chooses the adapter to use to obtain IP addresses through DHCP based on your selections in the Routing and Remote Access Server Setup Wizard. You can manually choose a LAN adapter from the **Adapters** list on the **IP** tab on the properties of a VPN server in the Routing and Remote Access snap-in.

If the static IP address pools are a range of IP addresses that are a subset of the range of IP addresses for the network to which the VPN server is attached, verify that the range of IP addresses in the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

- Verify that there are no packet filters on the profile properties of the remote access policy corresponding to VPN connections that are preventing the sending or receiving of traffic.

Unable to establish tunnel

- Verify that packet filtering on a router interface between the VPN client and the VPN server is not preventing the forwarding of tunneling protocol traffic. See Appendix A for information on the types of traffic that must be allowed for PPTP and L2TP/IPSec traffic.

  On a Windows 2000–based VPN server, IP packet filtering can be separately configured from the advanced TCP/IP properties and from the Routing and Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.

- Verify that the Winsock Proxy client is not currently running on the VPN client.
  When the Winsock Proxy client is active, Winsock API calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server.

  A proxy server–based computer allows an organization to access specific types of Internet resources (typically Web and FTP) without directly connecting that organization to the Internet. The organization can instead use private IP network IDs (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16).

  Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and a VPN server (for authorized Internet users) to facilitate both exchanges of information.

Appendix E: Deploying a Certificate Infrastructure

In a typical deployment, the certificate infrastructure is configured using single root CA in a three-level hierarchy consisting of root CA/intermediate CAs/issuing CAs. For VPN connections, issuing CAs are configured to issue computer certificates or user certificates. When the computer or user certificate is installed on the VPN client, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. When the computer certificate is installed on the authenticating server, the issuing CA certificate, intermediate CA certificates, and the root CA certificate are also installed. The issuing CA for the computer certificate installed on the authenticating server can be different than the issuing CA for the VPN client certificates. In this case, both the VPN client and the authenticating server computer have all the required certificates to perform certificate validation for both IPSec and EAP-TLS authentication.

When deploying a certificate infrastructure, use the following best practices:

- Plan your certificate infrastructure before deploying CAs.
- The root CA should be offline and its signing key should be secured by a Hardware Security Module (HSM) and kept in a vault to minimize potential for key compromise.
- Organizations should not issue certificates to users or computers directly from the root CA but rather should deploy the following:
    - An offline root CA
    - Offline intermediate CAs
    - Online issuing CAs

    This CA infrastructure provides flexibility and insulates the root CA and intermediate CAs from attempts to compromise its private key by malicious users. The offline root and intermediate CAs do not have to be Windows 2000 CAs. Issuing CAs can be subordinates of a third-party intermediate CA.
- Backing up the CA database, the CA certificate, and the CA keys is essential to protect against the loss of critical data. The CA should be backed up on a regular basis (daily, weekly, monthly) based on the number of certificates issued over the same interval. The more certificates issued, the more frequently you should back up the CA.
- You should review the concepts of security permissions and access control in Windows, because enterprise certification authorities issue certificates based on the security permissions of the certificate requester.

If you want to take advantage of auto-enrollment for computer certificates and the requesting of certificates using the Certificates snap-in, use Windows 2000 Certificate Services and create an enterprise CA at the issuer CA level.

For more information, see the topics titled "Checklist: Deploying certification authorities and PKI for an intranet" and "Checklist: Creating a certification hierarchy with an offline root certification authority" in Windows 2000 Server online Help.

Certificate revocation and EAP-TLS authentication

By default, the authenticating server checks for certificate revocation for all the certificates in the certificate chain sent by the VPN client during the EAP-TLS authentication process. If certificate revocation fails for any of the certificates in the chain, the connection is not authenticated and is denied. The certificate revocation check for a certificate can fail because of the following:

- **The certificate has been revoked.** The issuer of the certificate has explicitly revoked the certificate.
- **The certificate revocation list (CRL) for the certificate is not reachable or available.** CAs maintain CRLs and publish them to specific CRL distribution points. The CRL distribution points are included in the CRL Distribution Points property of the certificate. If the CRL distribution points cannot be contacted to check for certificate revocation, then the certificate revocation check fails. Additionally, if there are no CRL distribution points in the certificate, the authenticating server cannot verify that the certificate has not been revoked and the certificate revocation check fails.
- **The publisher of the CRL did not issue the certificate.** Included in the CRL is the publishing CA. If the publishing CA of the CRL does not match the issuing CA for the certificate for which certificate revocation is being checked, then the certificate revocation check fails.
- **The CRL is not current.** Each published CRL has a range of valid dates. If the CRL Next update date has passed, the CRL is considered invalid and the certificate revocation check fails. New CRLs should be published before the expiration date of the last published CRL.

This behavior can be modified using the following registry settings in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 on the authenticating server:

- **IgnoreNoRevocationCheck** When set to 1, the authenticating server allows EAP-TLS clients to connect even when it does not perform or cannot complete a revocation check of the client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate doesn't include CRL information. IgnoreNoRevocationCheck is set to 0 (disabled) by default. An EAP-TLS client cannot connect unless the authenticating server completes a revocation check of the client's certificate chain (including the root certificate) and verifies that none of the certificates have been revoked. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties.
- **IgnoreRevocationOffline** When set to 1, the authenticating server allows EAP-TLS clients to connect even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. The authenticating server does not allow clients to connect unless it can complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a server that stores a revocation list, EAP-TLS considers the certificate to have failed the revocation check. Setting IgnoreRevocationOffline to 1 prevents certificate validation failure because poor network conditions prevented their revocation check from completing successfully.
- **NoRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the wireless client's certificate. The revocation check verifies that the VPN client’s certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.

- **NoRootRevocationCheck** When set to 1, the authenticating server prevents EAP-TLS from performing a revocation check of the VPN client's root CA certificate. NoRootRevocationCheck is set to 0 by default. This entry only eliminates the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the VPN client's certificate chain. You can use this entry to authenticate clients when the certificate does not include CRL distribution points, such as those from third parties. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All of these registry settings must be added as a DWORD type and have the valid values of 0 or 1. The VPN client does not perform certificate revocation checking of the authenticating server's certificate and does not use these settings.

Because certificate revocation checking can prevent VPN access due to the inaccessibility or expiration of CRLs for each certificate in the certificate chain, design your certificate infrastructure for high availability of CRLs. For instance, configure multiple CRL distribution points for each CA in the certificate hierarchy and configure publication schedules that ensure that the most current CRL is always published and available.

Certificate revocation checking is only as accurate as the last published CRL. For example, if a certificate is revoked, by default the new CRL containing the newly revoked certificate is not automatically published. CRLs are typically published based on a configurable schedule. This means that the revoked certificate can still be used to authenticate because the published CRL is not current; it does not contain the revoked certificate and can therefore still be used to create wireless connections. To prevent this from occurring, the network administrator must manually publish the new CRL with the newly revoked certificate.

By default the authenticating server uses the CRL distribution points in the certificates. However, it is also possible to store a local copy of the CRL on the authenticating server. In this case, the local CRL is used during certificate revocation checking. If a new CRL is manually published to the Active Directory, the local CRL on the authenticating server is not updated. The local CRL is updated when it expires. This can create a situation whereby a certificate is revoked, the CRL is manually published, but the authenticating server still allows the connection because the local CRL has not yet been updated.

Using third-party CAs for EAP-TLS authentication

You can use third-party CAs to issue certificates for EAP-TLS authentication as long as the certificates installed can be validated and have the appropriate properties.

Certificates on the authenticating servers

For the computer certificates installed on the authenticating servers (either the VPN servers or the IAS servers), the following must be true:

- They must be installed in the Local Computer certificate store.
- They must have a corresponding private key.
- The cryptographic service provider for the certificates supports SChannel. If not, the certificate cannot be used and it is not selectable from the properties of the **Smart Card or Other Certificate** EAP type on the **Authentication** tab on the properties of a profile for a remote access policy.
- They must contain the Server Authentication certificate purpose (also known as an Enhanced Key Usage

[EKU]). An EKU is identified using an object identifier (OID). The OID for Server Authentication is "1.3.6.1.5.5.7.3.1".

- They must contain the fully qualified domain name (FQDN) of the computer account of the authenticating server in the Subject Alternative Name property of the certificate.

Additionally, the root CA certificates of the CAs that issued the VPN client user certificates must be installed in the Trusted Root Certification Authorities certificate store of the authenticating servers.

Certificates on VPN Client Computers

For the user certificates installed on VPN client computers, the following must be true:

- They must have a corresponding private key.
- They must contain the Client Authentication EKU (OID "1.3.6.1.5.5.7.3.2").
- They must be installed in the Current User certificate store.
- They must contain the universal principal name (UPN) of the user account in the Subject Alternative Name property of the certificate.

Additionally, the root CA certificates of the CAs that issued the IAS server computer certificates must be installed in the Trusted Root Certification Authorities store of the VPN client computers.

Summary

This paper described in detail the components and their associated design decisions for a Windows 2000-based remote access VPN deployment including VPN clients, Internet infrastructure, authentication protocols, VPN protocols, VPN servers, intranet infrastructure, AAA infrastructure, and certificate infrastructure. This paper also included detailed walkthroughs of PTPP and L2TP-based VPN deployments using computers running Windows 2000 and Window XP, details of firewall configuration, how to set up a VPN test lab, and a discussion of VPN troubleshooting tools and common VPN problems with suggested solutions.

Related Links

For more information about VPN, see the Windows 2000 Virtual Private Networks page at http://www.microsoft.com/vpn.

To download Microsoft L2TP/IPSec VPN Client, see Microsoft L2TP/IPSec VPN Client at http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/l2tpclient.asp.

For the latest information on Windows 2000, check out our World Wide Web site at http://www.microsoft.com/windows2000/default.asp.