



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

When Security Counts: Securing a Test Server with a VPN Connection

A virtual private network (VPN) is a logical link that connects shared or public networks. In this case study, a router-to-router VPN connection securely connects classroom computers to an isolated and secure internal network segment hosting an assessment server. This paper describes the design choices of a deployment for a router-to-router VPN connection using the Windows 2000 platform VPN server.

Copyright SANS Institute
Author Retains Full Rights



When Security Counts: Securing a Test Server with a VPN Connection
By Patricia Hulseay
January 13, 2003

Abstract

A virtual private network (VPN) is a logical link that connects shared or public networks. In this case study, a router-to-router VPN connection securely connects classroom computers to an isolated and secure internal network segment hosting an assessment server. This paper describes the design choices of a deployment for a router-to-router VPN connection using the Windows 2000 platform VPN server.

Background:

Our company, a large quasi-governmental organization, employs over 750,000 career and contract employees and provides customer services at approximately 38,000 retail outlets across the United States. This \$900 billion industry accounts for nearly 8% of the gross national product and is vital to business communications and commerce in the country. To maintain its current position in a technology driven market, the organization is working to optimize its network distribution systems.

Over the past decade as the organization has sought to increase revenue and maintain expenses, it has focused on automating its sorting and processing systems for network distribution. For example, a new software component built into the process system is able to read hand written addresses and print a bar coded zip code on each unit. Addresses written too poorly for the software to interpret are viewed through an online screening process by a person, who interprets the data and keys in the bar code for the correct zip code. This process permits each unit to be automatically sorted more efficiently and with fewer errors than can be done manually. No unit is ever handled again by a human until it is collected by the courier for delivery to the actual site.

The field automation has resulted in a significant reduction in semi-skilled force and the build up of highly skilled and technical support personnel to manage one of the largest networks in the world. The network that ties all of the retail facilities and corporate employees together is called the Internal Routed Network or IRN. The IRN also provides all communication services for its employees, including the training center instructors.

However, events of the past year have caused attention to be refocused on security, both physical and virtual. The primary organizational objective for the 2002 fiscal year is to strengthen the security of the IRN. The first step in the plan is to consolidate the delivery of all its internal network resources and services using a Microsoft environment. For example, there has been a centralization of the networking services, formerly supported at most of the 38,000 locations, into 15 regional centers around the country. The best illustration of this, perhaps, is the move from ccMail servers at each facility to Exchange mail server farms at these 15 regional centers.

In addition, all employees are receiving new computer systems. The new hardware will allow the organization to move from the Microsoft Windows 95 operating system to the Microsoft XP Enterprise edition operating system. One example of the Windows 95 operating system's lack of security is that the computers are not registered with a

domain controller. The Windows 95 operating system is now close to 10 years old and is a relatively insecure operating system when compared to the Windows XP, which offers a more robust security system. The new system is called the ACE for the Advanced Computing Environment.

The move to the ACE network is scheduled to occur over several months. This has become a critical issue at the central training facility where I am employed. Internal employees from around the country arrive with their own laptops and attempt to connect to the internal Windows 2000 network from our facility. Before ACE all organizational equipment was statically assigned an IP address in the private IP range of addresses by the local network administrator. When some of these “visiting” laptops connect to a different network, the local IT staff’s trouble shooting skills are really honed. Duplicate IP addresses can create havoc within the campus network. It also takes little imagination to understand the problems that can arise when inconsistent virus protection software or incompatible versions of application software are running on the network.

The new network security measures are important for the IRN and work well for most of the 38,000 facilities. The ACE program utilizes Active Directory (AD) services, Dynamic Domain Name Service (DDNS) and Dynamic Host Configuration Protocol (DHCP). These protocols have eliminated some of the connectivity issues mentioned in the earlier paragraph. The ability to push upgrades to software applications and to client software has minimized other issues.

But, the new security policy for desktop connectivity to the IRN, as implemented, has not worked well for the company’s in-house training facility. The instructors at this facility utilize many of the resources available over the Internet such as streaming video, webcasts, and online seminars. These resources have been particularly important not only for the continuing education of the training instructors but also for the class preparation. For example, several of the Microsoft instructors are working with the Microsoft Online Education program to prepare for certification exams as well as to support their in-house classes. In addition, there are a number of web seminars sponsored by numerous vendors about the new and emerging technologies, which are an important avenue for maintaining the expertise of our instructors.


The permissible use policy has become very important issue at the training facility. All employees and computers must authenticate in order to connect to the IRN for all network services from DHCP to Internet access. The new Usage Policy restricts many of the types of resources previously mentioned as important to the training instructor. The new firewall implementations, for example, prevent NetMeeting connections with computers outside the firewall which resulted in blocking the participation in a net meeting by an instructor.

The Network

Instructors at the training facility wanted continued access to those Internet resources to which they formerly had access through the IRN in order to maintain their certifications and support their classroom activities, while not abusing the desktop connectivity security policy for the IRN.

Initially stated, the goal is to design a network solution that will provide Internet connectivity to all of the Systems and Network Solutions classrooms which include the courses for Comptia, Microsoft, Cisco and SolarWinds training. The design moves the classroom and non-confidential communications off the IRN to the SansNet network.¹ The Network design to accomplish this goal is shown in Figure 1.

Figure 1 SansNet Network²



However, as a Microsoft Certified System Engineer and Cisco Certified Network Professional with ten years experience as a network administrator, this network design turned out to be relatively easy for me to implement. Equipment availability is not an issue because of the breadth of courses offered by the Systems and Network Solutions

¹ Note: I designed the SansNet network after my attendance at the Sans Conference in Washington, DC which is based on my understanding of defense in depth strategy for protecting networks connected to the Internet. It is possible for me to implement this strategy because I am providing non-essential services for classroom connectivity only. The IRN is the primary connection for instructors through their office desktop computers. The SansNet network is considered experimental and its secondary purpose is to permit instructors to experiment with new equipment, operating systems and applications to enhance their understanding of new technologies in a “real” world environment. Its primary purpose, of course, is classroom connectivity to the Internet.

² See Appendix A for a description of this graphic in more detail.

Team. Most of the routers and switches used in the SansNet network have been extricated from the various classrooms supported by the Systems and Network Solutions Team. With the upgrade of the computer hardware systems throughout the building to support the move to Windows XP, there is an abundance of older desktop computer systems that have been put to use on the SansNet network.

There is no wireless support on the IRN. Before the ACE program the wireless network was available only in the Net+ classroom and was used to support a teaching module. The wireless component is the only area of the design with which I was not familiar and, therefore, could not implement. The Network+ instructors established and are providing the technical support for this component of the SansNet network.

The Problem Statement:

As often happens, it is only after the network is designed and implemented that the real problem emerges. The new corporate desktop computer policy requires that only authenticated computers and users can access the IRN. Since the classroom computers use various operating systems depending upon the course being taught, the classroom computers are not registered to the organization's domain. In addition, the students attending these courses may or may not be employees of the organization and therefore may not have valid user accounts in the domain. Students are also not allowed to use a guest account to logon onto the IRN for security reasons. These factors all led to the implementation of the SansNet network.

The difficulty with the new network infrastructure arose when it was discovered that the students are often required to take a pre- and post-assessment test during the course. For internal students, these tests can be used to evaluate their career ladder promotions and salary considerations. Often the tests are used as a readiness measure for the student in his preparation for the real certification exam. In other cases, the external customer is interested in their employee's performance.

However, the assessments tests and the scored results database are maintained on a server on the IRN. This requires the student to logon into the test application on the IRN to take the test and to have the score recorded in the database. A report is generated on a weekly basis and posted to a web server where individual pass/fail scores can be accessed with the appropriate credentials from the Internet.

Simply stated, the objective is to allow all students from the classroom to logon onto an application located on a specific server on the IRN in order to take the pre- and post-assessment tests.

The Solutions Considered:

- One proposed solution involved relocating the assessment server to the SansNet Network but the staff who maintained the database opposed this solution.
- Another solution involved mirroring the site, while permitting the students access to the assessment application through the Internet and uploading the results to the internal server's database. The IT staff expressed an interest in finding a solution that did not add to their administrative responsibilities and this solution was discarded.

- The final approved solution involved establishing a gateway to gateway VPN from the classroom through the IRN connecting to the subnet containing only the test assessment application server.

By default, the Windows 2000 operating system supports only a single default gateway. When a Microsoft VPN client makes its connection to the VPN server, a new default gateway is set. The default gateway restricts access to any other segment of the network. Specifically stated, if a packet is destined for a network segment other than the far side of the tunnel, it is dropped. In this case study, the other side of the tunnel is a stub network which hosts only the assessment server. This prevents the classroom computer from being able to connect to the rest of the internal network for the duration of the VPN call. This is an especially important consideration because a VPN client has the same level of access any other internal host would have. This is a critical factor that resulted in the acceptance of the VPN as a solution.

By using a VPN solution, only those instructors who have appropriate credentials to logon to the classroom instructor computer can establish a VPN. This protects access to the resources of the assessment server. Those users who do not have the proper credentials cannot view the assessment server's network segment.

Rationale for Implementing a VPN Solution: ³

In many corporate networks, the departmental data, i.e. Human Resource data, experimental research data, and proprietary data, can be so sensitive that the department's network segment may be physically disconnected from the rest of the corporate network. While this protects the department's confidential information, information accessibility problems arise for those users who have access rights to this confidential data but who are not physically connected to the separate network segment. One solution to this problem is to create a VPN, which allows the department's network segment to be physically connected to the corporate network but separated by a VPN server.

The VPN:

There are many types of Virtual Private Networks (VPN). There are both hardware based (Cisco VPN 3000) and software based solutions (Microsoft Windows 2000 RRAS built-in VPN). For this case, I am using a software based solution, the Microsoft VPN.

A VPN is built by using tunneling protocols. The tunnel is created when one protocol is encapsulated within another protocol in order to cross a shared or public network. The tunnel is used to send data between two computers and works by emulating the properties of a point-to-point private link. Tunneling can be created either between source and destination routers or between two hosts. Tunneling can be point-to-point or point-to-multipoint. However point-to-point tunneling is much more scalable, and requires substantially less management overhead both from the establishment as well as maintenance stand point. Another advantage of a VPN is that it utilizes the existing infrastructure and can be rapidly implemented.

³<http://www.microsoft.com/windows2000/techinfo/howitworks/communications/remotefaccess/vpnoverview.asp>, page 3

Examples of tunneling protocols are General Routing Encapsulation (GRE), Layer 2 Forwarding protocol (L2F), Layer 2 Tunneling Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP) and Encapsulation Security Protocol (ESP).⁴ The Microsoft Windows 2000 Server operating system natively supports two protocols that will establish a VPN, the Layer 2 Tunneling Protocol (L2TP) and Point-to-Point Tunneling Protocol (PPTP). The PPTP protocol was selected for this solution because of ease of configuration. No Public Key Infrastructure (PKI) or certificate server is required with a PPTP implementation for authentication services on the VPN. Authentication is accomplished with user ID and password. PPTP automatically encapsulates and encrypts without the extra overhead that is required by L2TP for these services. IPsec protocol must be used in conjunction with L2TP for encryption.

Since Microsoft's implementation of PPTP uses GRE in its VPN solution, I will look briefly at how GRE works. Designated in 1994 in RFCs 1701⁵ and 1702⁶, GRE was unusual because it could encapsulate more than twenty (20) different types of protocols in its protocol field. Most tunneling protocols were protocol specific and capable of encapsulating only a single type of protocol. GRE tunneling incorporates three different protocols to encapsulate, provide carrier service and transport the packet.

A GRE tunnel is configured between the source (ingress) and the destination (egress) router. On a router-to-router VPN connection, the packets sent from either router across the VPN connection typically do not originate at the routers. The local router determines that the packets should be forwarded across the tunnel. The packet is encapsulated with a GRE header, transported across the tunnel to the tunnel end point address, where it is stripped of its GRE header so the original packet can be forwarded to its final destination. All of this is completely transparent to the two computers or users who are communicating with each other.

Considering a VPN Solution:

One key factor that contributed to the choice of a VPN solution is that any data behind the VPN server would have to be accessed through the VPN server.⁷ This means that for anyone to access the assessment application and/or the scored records in the database associated with the assessment tests, a VPN connection would be required.

Since the assessment server could only be accessed through the VPN and the instructor can control when the VPN service is running, this solution appeared to be able to protect the assessment server from unauthorized access even by the classroom computers. Inasmuch as the instructor computer is the VPN gateway (router) for the classroom computers, the first step is to disable any server services not warranted on this machine. The instructor machine's primary usage beyond providing tunnel ingress

4 Mason, Andrew G., Cisco Secure Internet Security Solutions, Cisco Press, page 368-369

5 Hanks, S., et al. "Generic Routing Encapsulation (GRE)", 1994, <ftp://ftp.rfc-editor.org/in-notes/rfc1701.txt>

6 Hanks, S., et al. "Generic Routing Encapsulation over IPv4 Networks", 1994, <ftp://ftp.rfc-editor.org/in-notes/rfc1702.txt>

7 <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windows2000serv/deploy/connfeat/vpnso1.asp>, page 1

for the student computers to the assessment server is to display the course presentation materials usually in a Power Point format.

Server Configuration:

The instructor machine runs the Windows 2000 Server operating system with Service Pak 2 installed as a member server. The computer is Compaq ENP III-800 MHz machine with 10 GB hard drive and 256 MB RAM. The hard drive is formatted using NTFS. Most services, such as IIS, and DNS, were not installed during the setup process.⁸

Based on recommendations found in the Microsoft Security Operations Guide for Windows 2000 Server in Chapter 4: "Securing Servers Based on Role", the member server baseline policy has been used to configure the Instructor server. Microsoft recommends the hisecws.inf security template for application servers. However, this security template requires the use of IPSec communication between clients and servers and between servers and servers. Since PPTP is the protocol planned to create the VPN, the hisecws.inf template is not an appropriate choice. Instead the securews.inf file is used as the policy template. The securews.inf template is a somewhat less restrictive level of security, included with both Windows 2000 server and Workstation installation software.

The member server baseline policy and the securews.inf policy⁹ are cumulative policies that address password complexity requirements and the account lockout policy settings for lockout reset and duration. These policies also incorporate the audit policy, which is set to track successful and failed application access. The security option restricts the display of the last username at logon. Access to the registry has been restricted to the administrators group by the registry access control settings. The file access control and service configuration settings can be restricted through this policy setting as well. No additional measures were taken.

The Design Considerations:

As previously noted, the VPN configuration is a gateway to gateway VPN solution. The instructor computer in the classroom served as the gateway (router) from the classroom to a server in the Information Technology (IT) center that served as the gateway (router) to the assessment application server's subnet.

Since we were implementing a VPN solution wholly contained within our own intranet environment we elected to use a PPTP connection for two reasons, ease of implementation and total cost of ownership. PPTP was designed to provide authenticated and encrypted communications between a client and a gateway or between two gateways, without requiring a public key infrastructure, reducing the total cost of ownership. The PPTP protocol is much simpler to implement and requires less administrative overhead because authentication is accomplished by user ID and password. PPTP protocol has some vulnerability associated with its use. However, since the VPN is totally contained on the intranet, there is no risk from the Internet and

⁸<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodtech/windows/windows2000/staysecure/default.asp> Appendix B Default Windows 2000 Services, pages 163-165

⁹ Shinder, Thomas w., Debra Littlejohn Shinder and D. Lynn White. *Configuring Windows 2000 Server Security*, Syngress Press, 2000, pages 144-146

minimal risk from internal network users.

Using the Point-to-Point Tunneling Protocol (PPTP), built on PPP¹⁰ (see RFC 2637) a TCP connection is established for tunnel maintenance and the Generic Routing Encapsulation (GRE) protocol encapsulates PPP frames for tunneled data. Once the encapsulated PPP frames reach their destination on the network, the frame is decapsulated and forwarded to its final destination on the protected network segment.

The VPN server or answering router is configured with persistent connections, which normally requires that the LAN interface remain permanently connected. But the VPN client or calling router can be disabled from the classroom by right clicking on the VPN connection in the RRAS MMC when the VPN connection is not utilized.

To prevent the calling router from making unnecessary connections, one can restrict the calling router by using Demand-dial filtering. Demand-dial filtering is configured to drop packets of the types of IP traffic that do not cause a demand-dial connection to be made and to pass the packets of the types of IP traffic that should be forwarded. Filtering on the VPN connection is configured on the inbound interface of the VPN server limiting the type of IP traffic that can cause a connection to be activated and providing another layer of security. One can configure demand-dial filtering by right-clicking the demand-dial interface in the Routing Interfaces node in the Routing and Remote Access MMC, and then clicking Set IP Demand-Dial Filters.

The next choice was whether to configure the VPN as a one way initiated or two way initiated demand dial connection. Since the assessment server would most likely never need to initiate a connection, the one way initiated demand dial connection was chosen. In a one-way initiated connection, one VPN router is always the calling router and one VPN router is the always the answering router.

One-way initiated connections require the following:

- The answering router is configured as a LAN and demand-dial router.
- A user account is added for the authentication credentials of the calling router that is accessed and validated by the answering router.


A demand-dial interface is configured at the answering router with the same name as the user account that is used by the calling router.

Configuring VPN Server:

Figure 2 is an illustration of the VPN design used for the classroom connectivity to the assessment server located on the IRN.

¹⁰ K. Hamzeh, et al., Point-to-Point Tunneling Protocol, 1999, <ftp://ftp.rfc-editor.org/in-notes/rfc2637.txt>, page 1

Figure 2 Classroom VPN



The classroom has a server running Windows 2000 Server with Service Pak 2 that acts as both remote access VPN client and demand-dial router. All computers in the classroom are connected to the 172.16.1.0/24 network (subnet mask 255.255.255.0). The classroom router (VPN Class) has a Spare interface that is assigned the private IP address 172.16.1.1.

The VPN server is a computer running Windows 2000 Server with Service Pak 2 that acts as both remote access server and demand-dial router. All computers in the assessment network are connected to the 10.0.0.0/24 network (subnet mask 255.255.255.0). The VPN router (VPN SRV IRN) has an Inside interface that is assigned the private IP address 10.0.0.1. All computers are in a workgroup.

To deploy the router-to-router VPN connection in this case, the following configurations are required:

- Configure and enable the Routing and Remote Access service on Classroom VPN.
- Configure a demand-dial interface on Classroom VPN.
- Configure and enable the Routing and Remote Access service on the VPN SRV IRN.
- Configure a demand-dial interface on the VPN SRV IRN.

- Establish the router-to-router VPN connection.
- Configure routing.

The step by step procedures for these configurations are not included here for brevity, since these procedures are graphically shown in most Windows 2000 Server books on the market.

Router-to-Router VPN Connections.¹¹

The calling router or the VPN client (VPN Class) initiates the connection. The answering router or the VPN server (VPN IRN) listens for connection attempts on predefined ports for the appropriate protocols. In this case, port settings are enabled on the IP filter to allow protocol 47 or GRE. Port settings are also enabled to permit source and destination ports 1723 for TCP. With these port filters enabled, the only type of connection allowed is the PPTP protocol. Once the VPN server receives a valid connection attempt from the calling router, it responds to the request to create a connection. The calling router authenticates itself to the answering router. When using a mutual authentication protocol such as MS-CHAP v2, the answering router also authenticates itself to the calling router.

The VPN routers that initiate or receive VPN-based demand-dial connections in this case study consist of the following components:

- **ROUTING AND REMOTE ACCESS SERVICE:** The Routing and Remote Access service must be configured as a VPN on both the calling and answering router server. This service is initiated using the Routing and Remote Access Server Setup Wizard. The wizard will start from the Administrative Tools menu by selecting RRAS. When the RRAS MMC opens, select the designated RAS server and choose Configure and Enable Routing and Remote Access.
- **PORTS:** Ports are software interfaces used as communication channels that support a single connection. Virtual private network (VPN) ports are logical ports. They are created when the VPN is configured through the RRAS wizard.
- **DEMAND-DIAL INTERFACES:** A demand-dial interface is configured on the calling router with the configuration information that will be used to complete the connection. This information includes the port type to use, the address used for the connection, which can be an IP address or domain name, the authentication methods, the encryption requirements, and authentication credentials. In this case study the port type uses TCP as the transport and GRE as the carrier. All computers are in a workgroup so an IP address rather than a domain name is used. Authentication methods are user ID and password.

The calling router is configured for a VPN connection in the usual manner to setup the one-way initiated connection. The answering router must have a static pool of addresses or DHCP configured to provide the calling router a valid IP address from the pool for that network. In this case the static pool is configured in RRAS for the server.

¹¹<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/network/deploy/depovg/vpnrote.asp>

The static pool of IP addresses configured for that server will be used by all dial-up connections. For example, if L2TP is also configured on the server as well as PPTP, the L2TP protocol would use the same static pool of IP addresses. The calling router must have an account specified that has permission to access the answering router. The answering router has either a user account with dial-in permissions, a remote access policy or will pull the account information from a domain controller if it is an active directory member. In this case, the answering router is a stand alone server with a user account configured to allow connection for the calling router. The account the calling router uses is also given a static route upon successful completion of the call to the answering router so that it can traverse the network on the other side of the answering router.

- **USER ACCOUNT:** To authenticate the calling router in a one-way initiated connection, the calling router's credentials must be the same as a corresponding user account in the answering router.
- **ROUTES:** The IP routes in the routing tables of the VPN routers must be configured to use the correct demand-dial interface to forward traffic across a router-to-router VPN connection. Static routes also must be configured on the answering router's user account in order for the return traffic to find the correct return path.
- **REMOTE ACCESS POLICY:** A separate remote access policy for the demand-dial connection is not required. Demand-dial access must be granted to the authenticating user account. The calling router can initiate a VPN connection based on an administrator action such as when the VPN connection is re-enabled. If the VPN is enabled, packets which match a route used by a VPN-based demand-dial interface will be forwarded. Once authenticated and authorized the VPN client acts as a router forwarding packets between nodes in the classroom and the answering router.

The answering router listens for VPN connection attempts. The answering router checks the credentials of the calling router's user account by comparing those credentials to its own user account credentials in order to authenticate and authorize the VPN connection before data is accepted. The answering router acts as a router passing packets between the assessment server and the calling router.

Testing the VPN:

In order to ensure the VPN solution meets the stated objective, the Sans Team instructors participated in a mock assessment test. IT Staff also attended the mock testing as observers. To demonstrate the ease with which a VPN could be established, the testing began with the calling router initiating a VPN connection with the answering router. The calling router connected without delay.

Instructors were provided with the directions to connect to the assessment application using a web browser. Once the instructors logged into the assessment application, they were able to answer exam questions without technical difficulty. The instructors reported on the follow up written evaluation that there appeared to be no significant difference in the test taking situation from the test-taker perspective. A final test demonstrated to the observers that instructors were unable to access the Internet from the classroom with the VPN enabled. Nor were the instructors able to access any other internal network

server.

The results of the mock exam were automatically forwarded to the database on the assessment server and a report was generated for posting to the web server. The IT Staff evaluated the report generated. The IT Staff concluded that the exam scores were correctly recorded in the database, resulting in an accurate report.

Conclusion:

One of the advantages of a VPN mentioned earlier in this paper is the fact that a VPN can be readily implemented. This design takes advantage of the instructor servers already integrated into the classroom setup and the existing intranet infrastructure. The PPTP protocol is relatively simple to implement and requires less administrative overhead than other tunneling protocols. No additional equipment is required for this implementation making this a cost effective solution.

The Microsoft VPN solution requires the configuration of a VPN connection and a demand-dial interface on the classroom server. To meet the IT staff requirement for a solution that did not add to their administrative responsibilities, the configuration of the IT Center server was kept to a minimum by using the one-way initiated VPN connection.

By using a VPN solution, only the instructors with the appropriate credentials can logon to the classroom instructor computer and establish a VPN to the protected assessment server. The time required by the instructor to re-enable the VPN connection is not an issue since it takes only a fraction of a minute. Students are only able to logon to the assessment application once the instructor has enabled the VPN connection. Then, students are able to connect to the assessment server by typing in the address of the assessment server in their browser. They are able to logon to the application and complete their assessment tests. The process appears to be completely transparent to the user.

After considerable discussion with IT staff, the VPN solution satisfies the IT staff requirements. The VPN solution also meets the objective, which is to allow students from the classroom to logon into an application located on a specific server on the IRN in order to take the pre- and post-assessment tests. Equally important, it is possible to accomplish these objectives with a minimum of expense to the organization.

© SANS Institute

On the inside interface the 3640 takes the inside or local address, which is not globally unique and translates that address to a globally unique address as required on the Internet.¹² This process is called NAT for Network Address Translation or NAT.

The next layer of security is the Pix 515 firewall or Pix1. Pix1 is configured to NAT. In this instance, Pix1 is taking an internal set of IP addresses (the 10.103.0.0/27 and NATs to another pool of internal IP addresses on the 10.104.0.0/27)

The first of two screened subnets or DMZ is on the internal interface of the Pix1 firewall. The first DMZ hosts all of the servers and services that must communicate with the outside network or the Internet. Some of these services must also be available to the internal network user. Those services, which must also be available to Internet users, are protected by static translations in Pix1.

The services planned are an exchange gateway, whose services are not currently operational. The long range plan is to install a Windows 2000 Exchange front end server to provide instructors with an internal team email system. Installing an Internet connection and maintaining the Microsoft Exchange 2000 Server will provide experience for those instructors who will be teaching the Microsoft Exchange 2000 Administration course to internal IT staff at the 15 regional centers. In this design the Exchange 2000 back end server is located on the internal classroom network.

The multimedia server hosts a number of video clips which are available to registered students for self study exam preparation. These clips are available from the classroom as well as across the Internet. Each class has a logon guest account with a six month time limitation set on the account.

The web server contains a mirror of the web pages on the corporate server which are specifically related to the SANS Team work product. The mirror web server enables the instructors to develop web pages and test the frames and links before they are placed on the corporate server.

Also hosted on this subnet is the syslog server for the network and the SolarWinds application, which performs some network monitoring functions. The external master DNS server for the SANS Team is on the DMZ. SolarWinds and Syslog are not accessible from access from the Internet.

The Pix2 firewall performs stateful packet filtering on the SansNet. Stateful packet filters listens to all communications, allowing only those packets expressly permitted by the Access Control List (ACL). Retaining those communications initiated from the internal network in memory, the stateful packet filter passes only those responses returned on expected ports from expected IP addresses. The role of this firewall is to protect the internal network from Internet Intruders. The Pix2 is configured to NAT from the 10.102.0.0/27 network to the 10.103.0.0/27 network.

The second DMZ is attached to Pix2. This DMZ will host a Remote Access Server (RAS) utilizing an AS5300 configured to authenticate to a Windows 2000 Active Directory (AD) located on the Domain Controller (DC) on the same network segment. The external slave DNS is located on this subnet. This DMZ is not currently activated.

¹² Akin, Thomas, Hardening Cisco Routers, O'Reilly, Sebastopol, CA., 2002, page 84-85

The Cisco 6509 core switch with Multilayer Switch Feature Card Module (MSFC) is configured with nine VLANs. The first VLAN is the management VLAN. The 6509 is configured to receive packets from the various VLANs and NAT the private 192.168.x.0/24 IP address to the private 10.102.0.0/27 pool of IP addresses. Each of the six classroom VLANs is switched from the classroom through a Cisco Catalyst 3524 switch, which is trunked to the 6509 Switch. The IP address for each classroom is on a private 192.168.x.0 network. The x in the third octet is consistent with the physical room number of the classroom for monitoring purposes.

The eighth VLAN is the wireless network. The wireless network is used exclusively by the instructors as an alternative to the IRN. Each instructor has a laptop containing a wireless network card for access to the SansNet. (Their desktop machine is connected to the IRN.)

The ninth VLAN is the internal server LAN. The internal server LAN hosts a variety of services that is used by the instructors. Included on this network are the Anti-Virus application used to push virus protection upgrades to instructor laptops and classroom computer equipment. Another important server is the image server, which contains the course images for each of the classroom servers and some of the desktops images. The Exchange 2000 Server, the internal DNS server and the DC round out the services found on the internal network. Additional servers and/or services may be added to the network as warranted either for instructor training initiatives or to support instructor work products.

© SANS Institute 2003, Author

References

- 2 Appendix A this paper for a description of this graphic in more detail.
- 3 Microsoft Corporation, A VPN Overview, currently online (01/10/03),
(URL)<http://www.microsoft.com/windows2000/techinfo/howitworks/communications/remotearchive/vpnoverview.asp>
- 4 Mason, Andrew G., Cisco Secure Internet Security Solutions, Cisco Press, 2001
- 5 Hanks, S., et al. "Generic Routing Encapsulation (GRE)", 1994,
(URL) <ftp://ftp.rfc-editor.org/in-notes/rfc1701.txt> (01/01/03)
- 6 Hanks, S., et al. "Generic Routing Encapsulation over IPv4 Networks", 1994, (URL)
<ftp://ftp.rfc-editor.org/in-notes/rfc1702.txt> (1/01/03)
- 7 Microsoft Corporation, A VPN Solution, currently online (01/10/03),
(URL)<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windows2000serv/depoy/feat/vpnso1.asp>
- 8 Microsoft Corporation, Microsoft Security Operations Guide for Windows 2000 Server, "Appendix B: Default Windows 2000 Services", currently online (01/10/03),
(URL)<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodtech/windows/windows2000/staysecure/default.asp>
- 9 Shinder, Thomas w., Debra Littlejohn Shinder and D. Lynn White. Configuring Windows 2000 Server Security, Syngress Press, 2000
- 10 Hamzeh, K., et al., Point-to-Point Tunneling Protocol, 1999,
(URL) <ftp://ftp.rfc-editor.org/in-notes/rfc2637.txt> (01/01/03)
- 11 Microsoft Corporation, TechNet article, currently online (01/10/03),
(URL)<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/network/depoy/depovg/vpnroute.asp>
- 12 Akin, Thomas, Hardening Cisco Routers, O'Reilly, Sebastopol, CA, 2002



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

Hong Kong Advanced Forensics Seminar	Hong Kong, Hong Kong	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Sydney 2009	Sydney, Australia	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Vancouver 2009	Vancouver,	Nov 14, 2009 - Nov 19, 2009	Live Event
SecurityByte 2009	New Delhi, India	Nov 17, 2009 - Nov 20, 2009	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	Geneva, Switzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS San Francisco 2009	OnlineCA	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced