Global Information Assurance Certification Paper

How To Secure Your Small To Medium Size Microsoft Based Network: A Generic Case Study

GSEC Practical Assignment V1.4b Option 1

By
Jerry L. Goodman
July 11, 2003

Introduction:

If you run a business in today's world undoubtedly you also run a network of computers, printers, and other devices as well. Though the network may not be the primary focus of your business it probably plays a key role in how it functions. That being said it is obviously just as important to protect your network from theft and damage as it is to protect the other property and assets of your business. In this paper I intend to explain the basic process of securing a small to medium sized network. I will create a make believe company network and give some examples of how to secure the network with some commonly used products and techniques in a case study format. I will use references to freely available information on the Internet to help me secure this network. The examples are based on Microsoft and Cisco platforms because they are the most commonly used platforms and they are the ones I am most familiar. I have implemented, evaluated, or researched each solution I recommend either for customers or my companies network.

Security for even a small network is not a simple task. This document just scratches the surface. The main function of this paper is to cover the main things you need to do. It is not meant to cover everything in detail but to give you a way to get started with pointers to locations for more information.

Securing your network is a little different than securing other types of assets but the overall idea is the same. There are four questions to answer, "What do I want to protect?" "How do I protect it?" How do I know if the protection is defeated?" and finally "What do I do if the protection is defeated?"

We will also incorporate the use of a layered defense or defense in depth.

Defense in Depth:

Defense in depth is a concept that comes to us from the military and goes far back into history. The concept is a layered defense like a castle or fortress. You have outer walls surrounded by moat with a large open field in front of it. The walls are a pretty good defense but they can be climbed or broken down. The field creates an area for your archers to fire on an enemy as they approach. That makes it harder to move things in to break the walls like catapults. The moat makes it harder to place ladders to climb the wall. All of those defenses serve to slow the enemy down while you ultimately protect everything inside the castle. A more recent example is in the movie "Die Hard". The criminals must break through a series of seven locks to get to the treasure. Each lock requires a different skill set and tools to defeat. The final defense is Bruce Willis and the Police Department. These are examples of Defense in Depth; multiple layers that require time and skill to defeat even if they are small inexpensive layers.

What do I want to protect?

Inventory:

The first question you'll need to answer is, "What do I want to Protect?"
That question is answered by taking an inventory of everything that is connected to your network. You need an inventory to tell you what you need to protect. Your assets include computers, network devices, software, and data to name a few.

One of the primary tools for maintaining, troubleshooting, and securing your network is good documentation. An inventory allows you to list your assets and determine the threats and vulnerabilities to those assets. Thus answering the question of "What do I need to protect?"

For our example we'll use a small company called Goodman Inc.
Goodman Inc. is the maker of specialized hardware and software tools that are sold directly by the company through a mail and phone order catalog and via a web server. The company is head quartered in Kansas City with offices in Wichita and Oklahoma City.
Kansas City uses a T1 line to connect to Wichita and another to connect to the Internet. Oklahoma City uses a T1 connection to a local ISP and utilizes a LAN-to-LAN VPN to connect to Kansas City. Remote users connect via VPN to Kansas City or by Dialup if Internet access is not available.

The assets break down as follows. There are 150 users in Kansas City 100 desktops systems 50 laptops, 25 systems in Wichita 10 desktops and 15 laptops, 10 laptops in Oklahoma City. All workstations are running Windows 2000 Professional. There are a total of 13 Intel based servers, 11 in Kansas City, 1 in Wichita and 1 in Oklahoma City. There are 5 Cisco Routers 3 in KC and 1 each in Oklahoma City and Wichita. There are two PIX firewalls one in KC, one in Oklahoma City, as well as a remote access server (RAS) in KC for dialup access.

Servers:

1. An external web server that uses an internal SQL database to provide sales over the Internet.
2. There is an internal web server that provides the companies Intranet.
3. An external SMTP server that is the SMTP gateway
4. An external Web-mail server that provides web access to e-mail
5. An internal e-mail server
6. A domain controller (DC)/ certificate server (CA)
7. A remote access server for dialup
8. An SQL server
9. A backup /File server (DC)
10. An RSA SecureID server for Strong authentication.
11. One DC in Wichita that does backup and file services.
12. One DC in Oklahoma City that does backup and file services.

13. Company data including customer, HR, sales and inventory information.
14. The company network including routers, firewalls, switches, etc…

Figure 1: Goodman Inc.

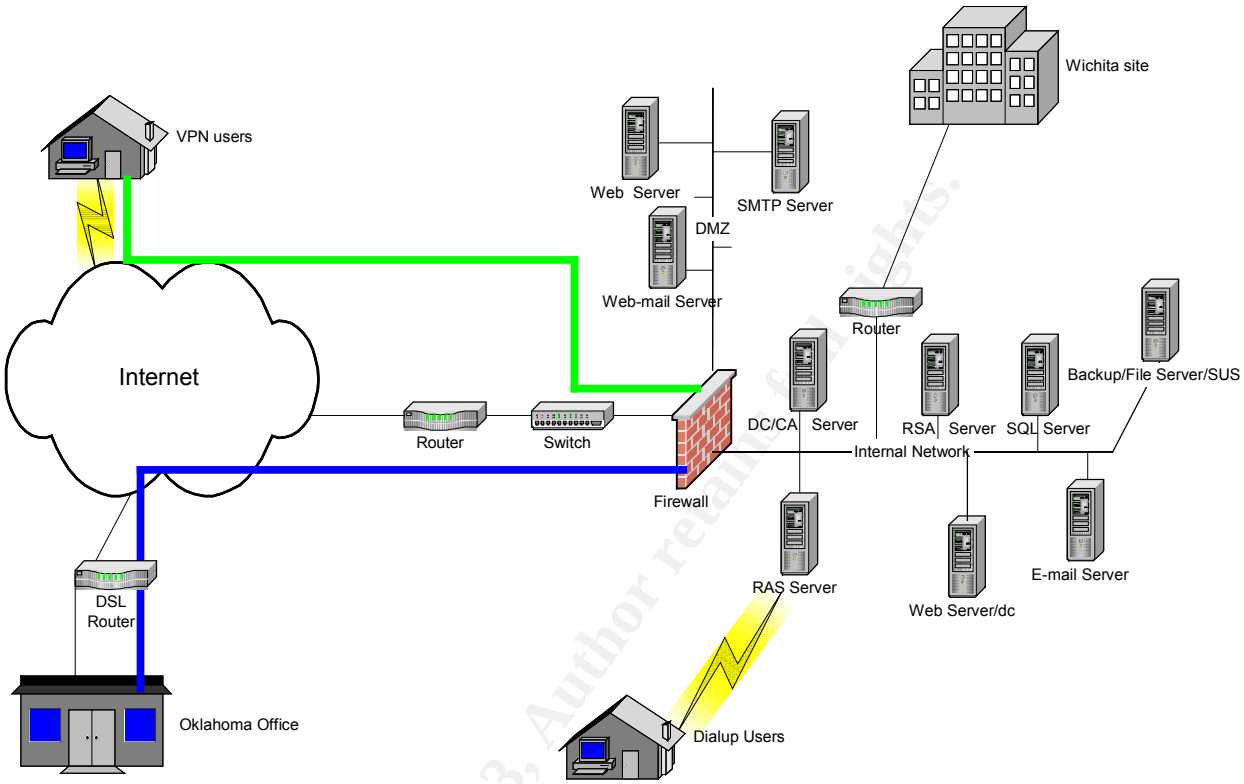

Risk Assessment:

If you think of network security as an insurance policy then you need to determine what to insure, and then prioritize those things. Prioritize by what it would cost to replace or fix and the cost to your business if it were unavailable. After you prioritize your inventory list you'll need to decide what threats you really need to protect them from. For example, if you were insuring items in your home that you value your two most valued possessions might be a $10,000 Flat-Screen Plasma HDTV and a $10,000 Jacuzzi. Both are worth $10,000 however the likely hood that your Jacuzzi would be stolen is considerably less than that of the TV. Since the risk of theft is lower for the Jacuzzi, then you would not want spend as much to insure it against theft as you would for the TV.
The decision process for determining this is called Risk Analysis and Risk Assessment There are two methods of risk assessment Qualitative and Quantitative.
Quantitative risk assessment is the process of assigning a value to represent the likelihood that a type of event or threat will occur and the cost if it did.

For example you could set your risk scale from 0 to 100 with 0 meaning there is no threat that the event could happen and 100 meaning that it would positively happen. Then you could multiply the potential loss by that value to determine the risk. That means that if the threat factor were equal to 2 and the cost were 50,000 then the risk would equal 100,000. This gives you a rough gauge to use in prioritizing where to spend your efforts. It is however highly susceptible to error since it is difficult to accurately create the risk scale.
There is a good paper on Quantitative risk analysis with a table of risk factors at the Computer Security Resource Center (CSRC) The URL is:
http://csrc.nist.gov/nissc/1999/proceeding/papers/p28.pdf

Qualitative risk assessment is the process of determining the possible threats to your assets, their vulnerability to those threats and how critical the assets are to your operation. This is the most commonly used method and the method we will use in our example.

Once you have listed your assets, the possible threats, and the vulnerability to those threats, you will need to prioritize them by their criticality to your operation and then determine a defense against those threats.
Choosing what defense mechanisms to use is a process of deciding what level of risk you can afford to live with vs. the cost of that defense and the level of difficulty to install and maintain it.

More detailed information on Qualitative Risk Analysis can be found in the paper written by Tim Bass and Roger Robichaux titled **Defense-In-Depth Revisited: Qualitative Risk Analysis Methodology For Complex Network-Centric Operations.** The URL is
http://www.silkroad.com/papers/pdf/archives/defense-in-depth-revisited-original.pdf

Examples of some possible threats:

- Loss of Power
- Loss of Communication
- Accidents
- Loss of Data Integrity
- Computer Virus
- Abuse of privilege by employees
- Natural disasters
- Unauthorized Access
- Theft or destruction of computing resource
- Destruction of data
- Denial of Service attack

Figure 2: Example of Risk Analysis for Servers

Threat	Vulnerability	Damage
Loss of Power	High	Loss of access to System, Possible loss of Data
Loss of Communication	High	Loss of access to System
Accidents	High	Loss of access to System, Loss of Data, Loss of Hardware
Computer Virus	High	Loss of access to System, Loss of Data
Abuse of privilege by employees	High	Loss of access to System, Loss of Data, Loss of Hardware
Natural disasters	Low	Loss of access to System, Loss of Data, Loss of Hardware
Unauthorized Access	High	Loss of access to System, Loss of Data
Theft or destruction of computing resource	Medium	Loss of access to System, Loss of Data, Loss of Hardware
Destruction of data	High	Loss of access to System, Loss of Data
Denial of Service attack	High	Loss of access to System

We will create risk analysis documents for all of the items on our inventory list. Some of the items can be combined as we did with the risk analysis for the servers. The idea is to get a good idea of the things we need to protect and the threats we need to protect them from.

How do I protect it?

Protection:

The next question we need to answer is "How do I protect my assets?"
Once we have listed our assets and listed the threats to those assets and determined our risk we are ready to move on to prevention. We will now establish policies or rules to define the level of protection that we will require. Once we establish the policy we can then determine the methods we will employ to meet the policy requirements.

Policies:

Why policies? Policies are guidelines and rules that allow you to set protection levels and the punitive actions that can be taken for violations.
At first glance creating security policies sounds like a daunting task. Fortunately we have the benefit of being able to start with some security policy templates created by security professionals and made available by the SANS organization.
http://www.sans.org/resources/policies/
At Goodman Inc we will make use of these policies. They may not exactly fit our needs but we can modify them accordingly. The specific policies chosen and their content will be driven by the findings of the risk analysis. For example in our risk analysis of our servers power loss is determined to be a high risk for denial of service and a medium risk for loss of data. To address this our server policy requires redundant power supplies and battery backup. Our Backup and Recovery policy protects us from loss of data.

Let's look at one complete policy in detail then we will talk about some specifics of the other policies. We will start with the Server Security Policy Template from the SANS website. http://www.sans.org/resources/policies/
We took the policy template and made very few changes to create our policy. (See Appendix A.)
In the Policy we define some rules to be followed. First we stipulate that all servers must have configuration guides that are approved by the security department. This gives us the assurance that our servers meet our security requirements from the beginning.
We require all servers to be registered with our corporate enterprise management system. This management system is just a small database on our SQL server. This database contains information on all of our servers and computers. The information includes

- Make and model of hardware and devices.
- Software installed with versions and License info.
- Applications
- Contact information (internal and vendor)
- Major Patch and Service Pack revision.

Given the size of our company the information could be just as easily stored in an Excel spreadsheet on our file server.
We require the information to be kept up to date. We require a quarterly review by the IS department of this information for accuracy. We require that the Operating systems and software programs stay as up to date as possible. For this purpose we have a test/development server. This server runs Windows 2000 server, IIS, Exchange 2000, and

SQL. It contains copies of the internal and external web services and applications as well as small versions of the company e-mail and databases. We use this server to develop and test changes to the sales application as well as testing OS upgrades and patches. All Service packs and patches are tested on this system unless an emergency exception is approved by the security department. Our guideline on Hotfixes is that they be installed only as specifically required to address a specific problem. Service Packs will be installed when they are tested and approved.
We require that Operating System configuration should be in accordance with approved security department guidelines. To this end we have created a baseline security model for our servers and desktops.

Server Guidelines:

For Windows 2000 servers the guidelines are as follows.

- All servers are to be in a locked room with controlled access.
- All disk partitions are formatted NTFS.
- All passwords must adhere to the Password policy.
- All servers must be backed up daily. Incremental, differential, or full backup can be used for daily backups. Full backups are required weekly with a second full copy backup created monthly for archival.
- All backup media will be stored offsite on the first business day following the backup.
- Weekly and daily backup media can be overwritten on a bi-monthly basis.
- All Microsoft servers must have ERDs updated monthly and stored with Backups.
- All unnecessary protocols must be removed.
- All unnecessary services must be removed.
- Disable or delete all unused accounts.
- Disable the guest account.
- Authentication must be NTLMv2 or Kerberos no Basic or Anonymous.
- Disable caching of logon information.
- Set the paging file to be cleared at system shutdown.
- Restrict interactive login to Administrators only no anonymous logon.
- Restrict Floppy and CD-ROM access to interactive user only.
- No remote access to the registry.
- All servers must use company logon warning banner.

There are additional guidelines for specific servers such as web servers, file and print servers, Domain controllers, etc… Microsoft has recommendations regarding these guidelines. They are part of a Securing Windows 2000 Server guide at http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodtech/Windows/SecWin2k/06basewn.asp

They have a chapter (Chapter 7) on security hardening of servers with specific roles. They employ tactics such as running IISLockDown and URLScan tools for web servers as well as a specific baseline security model for Domain controllers.

We require that all servers log security information and retain it for review. We also require events to be reported to the security department as follows.

Security-related events will be reported to the security department, who will review logs and report incidents to IT management. Corrective measures will be prescribed as needed. Security-related events include, but are not limited to:

- Port-scan attacks
- Evidence of unauthorized access to privileged accounts
- Unusual events that are not related to specific applications on the host.

This type of policy creates a checks and balances system that allows the security department to help the IT department protect its assets.
There is also an enforcement clause that gives the policy teeth in case of willful or negligent violation of the policy. The policies define the requirements but in themselves they don't protect anything. For example traffic laws like policies define requirements. But the laws themselves don't stop people from speeding or going the wrong way down one-way streets. To do that you need to define consequences. That is why all of the policies should have an enforcement statement Example: *Any employee found to have violated this policy may be subject to disciplinary action, up to and including termination of employment.*

This gives you the idea of how I developed a policy by modifying a template from the SANS website and made it applicable to this environment. This should give you an understanding of how to set up policies and how they are used. The process for defining the others is the same just choose the policies needed to protect the assets that you have. Now I'll mention some of the specifics of some of the other policies.

- Acceptable Use Policy – This just defines what is and isn't acceptable usage of the company network resources. For example no downloading of pornographic material.
- Backup/Recovery Policy – This defines what is acceptable in terms of documentation of backup/restore process, as well as the storing of and length of time to keep backups. In our example all Weekly Full Backups are required to be kept for at least one month before overwriting and all Monthly Full Backups must be kept for at least twelve months. This allows us to go back for a whole year if necessary to find data in our monthly backups and working in conjunction with an annual archival process allows us to go as far back as necessary.
- Firewall Policy - Describes what is allowed to pass through the firewall by default and the process for acquiring exceptions to the policy. For example by default no traffic is allowed to pass through the firewall from the outside directly through to the inside with an Internet routable address. Only reserved addresses defined by RFC 1597 in the DMZ (Defined by the DMZ policy) are allowed to pass through the firewall.

- Anti-virus process – Defines the requirement for anti-virus software in the environment

Policy Solution:

Below is the list of the policies that we will implement.

- **Acceptable Use Policy**
- **Firewall Policy**
- **Anti-Virus Process**
- **Confidential Information Policy**
- **Digital/Analog Communication Line Policy**
- **Audit Policy**
- **Authentication Policy**
- **Intrusion Detection Policy**
- **DMZ Security Policy**
- **Password Protection Policy**
- **Remote/ Dial-in Access Policy**
- **Network Switch and Router Security Policy**
- **Server Security Policy**
- **VPN Security Policy**
- **Security Awareness Policy**

Methods of Protection:

To further protect against people breaking these policies we need methods to protect our network assets. Some examples are Antivirus Software, Firewalls, Encryption, Security Awareness training, Strong or two-factor authentication, and doors with locks.

Firewall:

Lets start from the outside and work our way in. Our first line of defense is the company firewall. First off, what is a firewall? A firewall is a computer or router that acts like a gateway between your internal network and the Internet or any other network connected to your network. A firewall's job is to block undesired access from other networks into yours while allowing the desired traffic.
It is an entity that enforces security policy on your boundary to the Internet.
What is called a firewall today is usually a combination of tools placed on a common device. Each of these tools has a specific role and supplies a critical part of your security defense.

Filters:

One of the jobs a firewall performs is filtering information from the outside to the inside based on a set of rules which you define. There are three types of filters, Packet filters, Proxies, and Stateful inspection filters.

Packet Filters:

A packet filter runs at the network layer of IP and examines incoming network traffic, compares each packet to a set of rules allowing only the traffic that matches the filter to pass. Packet filters are fast and they can strengthen security but they can be fooled. Packet filters rely on a combination of ports and addresses to validate traffic but they do not inspect the data. For example you may have an e-mail server inside your network that you wish to allow incoming e-mail traffic to be forwarded to. A packet filter could examine the traffic and allow traffic destined for IP address 10.100.100.1 port 25 to pass. It cannot however tell what type of content is really in the traffic. It also has to learn everything it needs to validate the traffic in each individual packet. It has no way to know what a conversation is much less if it has been hijacked or not. Let's take an example, suppose you are an officer in the army and your job is to screen mail to and from soldiers to make sure that they are not sending information regarding troop locations to the enemy. If you are functioning in packet filter mode then all you could do is check the addresses of the envelopes to make sure that none of the mail is addressed to or from the enemy. Now while you definitely want to stop any soldiers from sending or receiving messages directly to or from the enemy it certainly does not guarantee the security of the information.

Proxies:

That brings us to another type of filter, a proxy. A proxy runs at the application layer. It acts like a man in the middle by taking outbound requests from clients to remote servers as if the connection was meant for itself. It then establishes a connection to the remote server on behalf of the client maintaining 2 connections for the duration of the transaction. That means if I'm the firewall and I'm passing traffic to 10.100.100.1 port

25, I form a connection with the incoming traffic on the outside, then I build a connection on the inside and only pass SMTP. In our example of the soldiers' mail, that means that you would open every letter, read it, rewrite it, and send it on as if was from you. This is the most secure form of firewall but also has the lowest performance.

Stateful Inspection Filter:

So the problem is how do you have a secure firewall that doesn't kill performance? The answer is a stateful inspection filter. It checks the packets for certain kinds of states that the conversation would be in. For example FTP would have Gets, and Puts. It's kind of like spot-checking to make sure that you are actually running the protocol that you are supposed to. In our example of the soldiers' mail, that means that you would be checking the mail for phrases that a soldier might include in a letter back home as opposed to latitudes, longitudes, troop strengths and so on. This spot-checking process is not as time consuming as a proxy but it is a lot more secure than a packet filter.

As per our firewall policy no traffic originated out side the firewall is allowed to pass directly through to the internal network. It must first pass through the DMZ. DMZ stands for De-Militarized Zone. That is a military term describing a neutral area between two opposing forces. It doesn't mean that the neutral area is safe, just that it is not in the other forces area.

Figure 3: DMZ

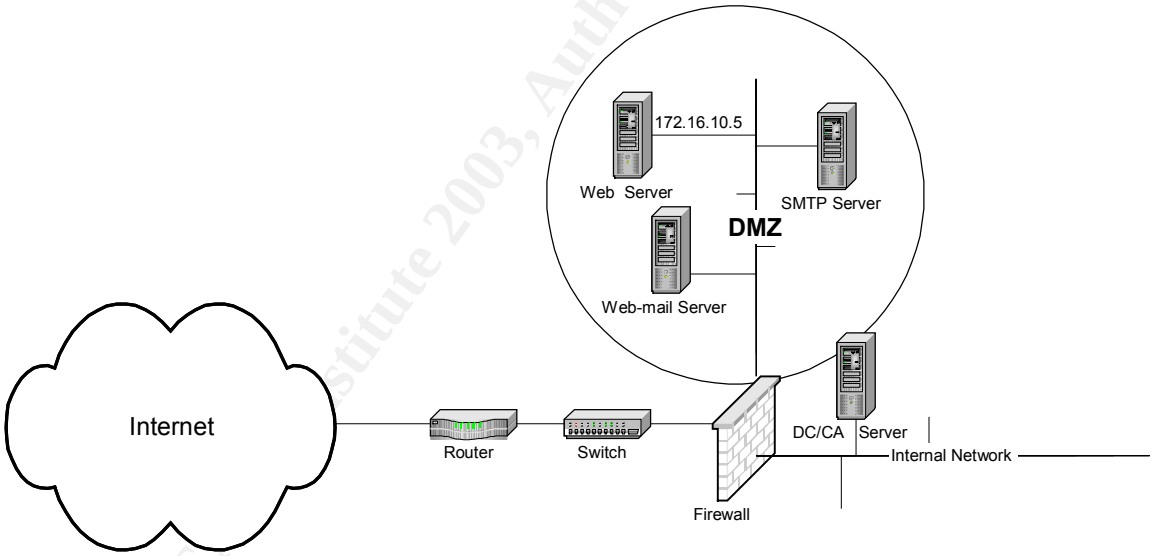

Let's take a second and discuss the public and private address ranges. If you understand IP addressing then skip the Public/Private Address range information and pickup at the Network Address Translation.

Public Address Range:

In the public address range there are 5 classes of IP addresses A, B, C, D, and E.
A, B, C, and D are the general address classes we would communicate on with D being reserved for multicast traffic and E reserved for experimental of future use. D and E classes are beyond the scope of this paper and will not be discussed here. Classes A, B,

and C are defined by RFC 796 a copy of which can be found at the following website. http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc0796.html

Basically the classes are a way to group IP addresses in such a way as to predetermine the ratio of possible hosts to possible networks for a given address. An IP address is 32 bits long and consists of four octets separated by periods in dotted decimal notation. 172.16.10.5 is in dotted decimal notation. Each octet is a decimal representation of a binary number. Each bit position of the octet is another power of two. The value of each bit is from left to right 128 64 32 16 8 4 2 1.

In our example 172 would = bit position 128 + 32 + 8 + 4 or 10101100. The entire address would be 10101100.00010000.00001010.00000101. Each IP address needs a mask to determine the host portion of the address from the network portion. In this case 255.255.0.0 is the mask.

In binary that mask would be 11111111.11111111.00000000.00000000. If you were to over lay the mask on the IP address, the network portion of the address is within the bits that are masked with a 1. The host portion is within the bits that are marked with a 0. So in this case the network address would equal 172.16 and the host address is 10.5 on that network.

Within each Class of addresses A, B, and C you have a different predetermined mask depending on the left most octet. Class A addresses always start with a 0 in the leftmost bit making the highest number that can be achieved in the first octet 127. Therefore class A consists of IP addresses that start at 1.0.0.0 and go thru 127.255.255.255. All class A addresses have a mask of 255.0.0.0. That means that any full class A address that begins with 1 – 127 in the first octet can have up to 16,777,216 hosts in each of 127 networks (Note: address 127 is a special address used for loop back and won't be discussed here).

Class B addresses begin with a 10 in the two leftmost bits starting the range at 128 – 191 in the first octet. They have a mask of 255.255.0.0, which gives them 16,384 networks and 65,534 hosts per network.

Class C addresses begin with 110 in the 3 leftmost bits starting the range at 192 – 223 with a mask of 255.255.255.0 creating 2,097,152 networks with 254 nodes per network. Note: reference Mastering TCP/IP for NT Server (Mark Minasi, Todd Lammle, and Monica Lammle.) Sybase books.

Private Address Range:

Now that you understand how IP addresses are divided into classes we are going to talk about the private address range.

The private address range is defined by RFC 1918 a copy of which can be found at the following website. http://rfc.sunsite.dk/rfc/rfc1918.html

Basically what it says is that for each class of address there will be a range of addresses that is not routed on the Internet. Those addresses are as follows.

Class A = 10.0.0.0 -thru - 10.255.255.255
Class B = 172.16.0.0 -thru – 172.31.255.255
Class C = 192.168.0.0 -thru – 192.168.255.255

This means that if you use any of these addresses on your internal network that no one can route traffic directly to those addresses. Couple this with NAT and you get a powerful layer of protection with your firewall.

Network Address Translation:

That brings us to another tool of most firewalls; Network Address Translation or NAT. NAT is defined by RFC 1631 a copy of which can be found at the following website. http://www.faqs.org/rfcs/rfc1631.html.

NAT is a way insulating the network addresses used inside our network from the outside. How that is done is by creating two sets of addresses, one on the outside of the firewall and one set on the inside. So if the external address of our web server were 10.12.11.12 then the NAT function of the firewall would translate that to the DMZ address of 172.16.10.5, which is the address that is assigned to the web server.

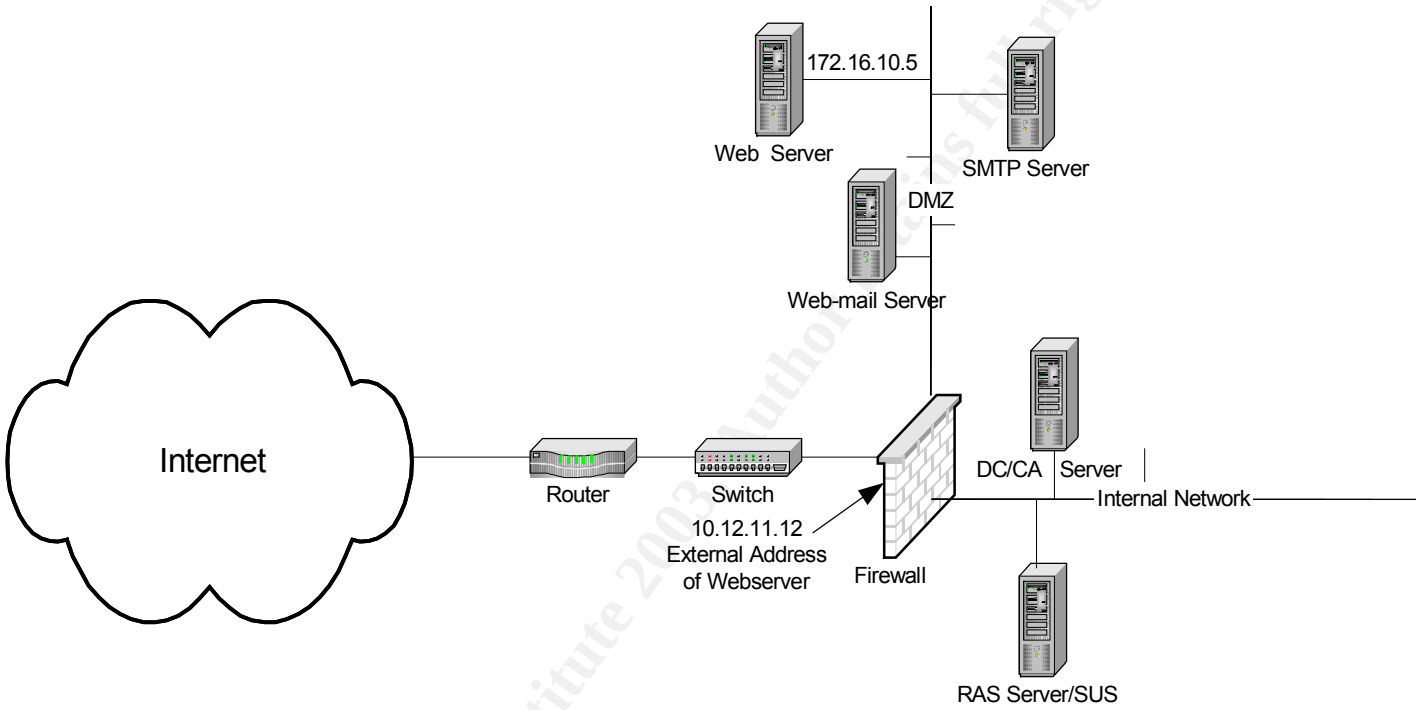

External Address	Internal Address	Device
10.12.11.12	172.16.10.5	Web server
10.12.11.2	172.16.10.1	PIX Firewall
10.12.11.6 thru 10.12.11.12	172.16.x.x	Outgoing connections

At first glance translating the address may not seem like such a great deal but let's take a closer look. With out address translation you could only use each IP address once and would soon run out of addresses. Having an address translation allows you to hide the internal address from the public. Also if you use an internal address that cannot be routed on the internet you place another road block in the way of possible attackers. Look at the external address in the example 10.12.11.12 is pretending to be a *real* Internet address in the public address range. While the address of 172.16.10.5 in the DMZ, is in the private address range.

Here is how it works. If you remember, we stated in our firewall policy that there

will be no traffic originating from the outside network (Internet) that will pass directly into the internal network. We accomplish this by utilizing the DMZ.

For example e-mail traffic that is destined for our internal email server must first pass to our SMTP relay operating in the DMZ. Our firewall has a rule that passes only SMTP traffic destined for 10.12.11.12 through to the NAT address of 172.16.10.5 on the DMZ. There the email is relayed to the internal network thru the firewall, which has a rule allowing SMTP traffic to pass only if it is from DMZ address 172.16.10.5 and destined for 192.168.10.23.

This method makes it harder for someone from the outside to attack the internal mail server. If the SMTP gateway is attacked our internal mail might not be directly affected. We might not be able to receive mail from the Internet while the SMTP gateway is down but we could still send mail out. We could send and receive mail internally and any messages that we already received would not have to be recovered. The SMTP gateway holds the mail for a very short period of time and forwards it to the internal mail server. Basically as soon as we correct the problem with the SMTP server we could be back in business. Whether it's a temporary denial of service attack or a complete corruption of the gateway, we restore the gateway and continue. No complicated data recovery should be necessary. It is possible but more difficult that an attacker could hack into the SMTP server but it is in the DMZ with only SMTP access to the e-mail server only. This is far better than having hacked into the internal e-mail server. Remember Defense in Depth consists of layers.

We do a similar thing with our web server access to the internal SQL server. We allow only specific application traffic to pass from DMZ to Internal network using only private addresses. We also require by the DMZ policy that all servers in the DMZ adhere to a hardened configuration guideline to further protect against hack

Figure 4: Address Translation

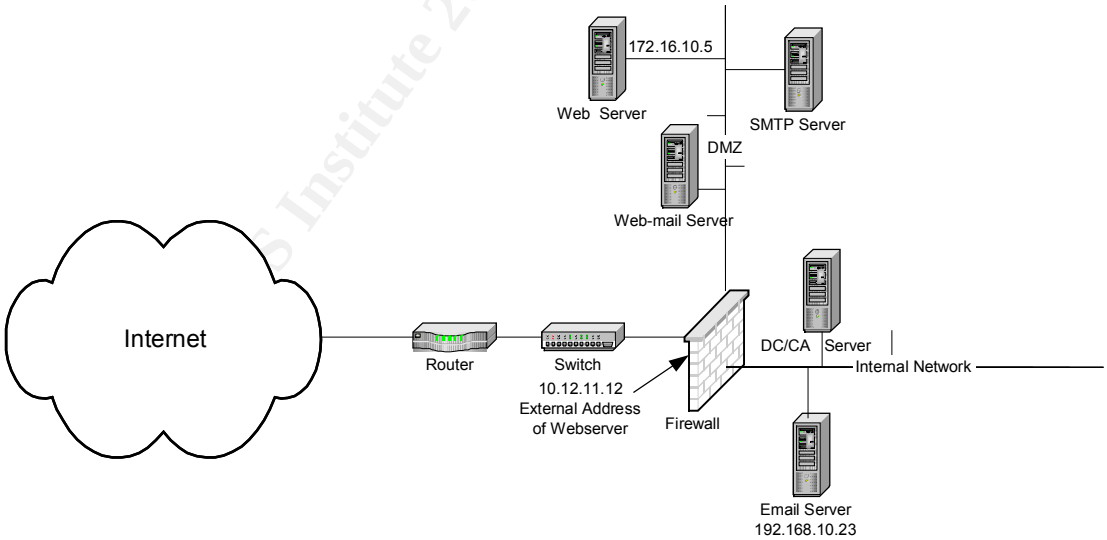

External Address	Internal or DMZ Address	Device
10.12.11.12	172.16.10.5	Web server
10.12.11.2	172.16.10.1	PIX Firewall
	192.168.10.23	E-mail Server
10.12.11.6 thru 10.12.11.12	172.16.x.x	Outgoing connections

Dynamic NAT:

That's how we control access originating from the outside but what about incoming traffic that was requested from the inside? In the previous example external address 10.12.11.12 always resolved to address 172.16.10.5. This one to one ratio is called a static NAT and while in works for the case above it does have some drawbacks. If we static map every address then we lose the protection of being able to hide our IP addresses from the outside. It would also mean that we would need a lot of external addresses to service our network.

Another way to use NAT is to dynamically translate the addresses. This is done by letting devices on the inside share a single external address or a pool of external addresses. This is sometimes called PAT or Port Address Translation.

How it works is simple. The firewall receives a request on the inside for a resource on the outside. The firewall creates an entry in its translation table for the request noting the destination IP address and port number as well as the source IP address port number. The firewall then forwards the packet to the destination address and port using one of its external addresses and a port number that it assigns as the source address. See step 1.

Figure 5: Address Translation Step 1

Step 1.

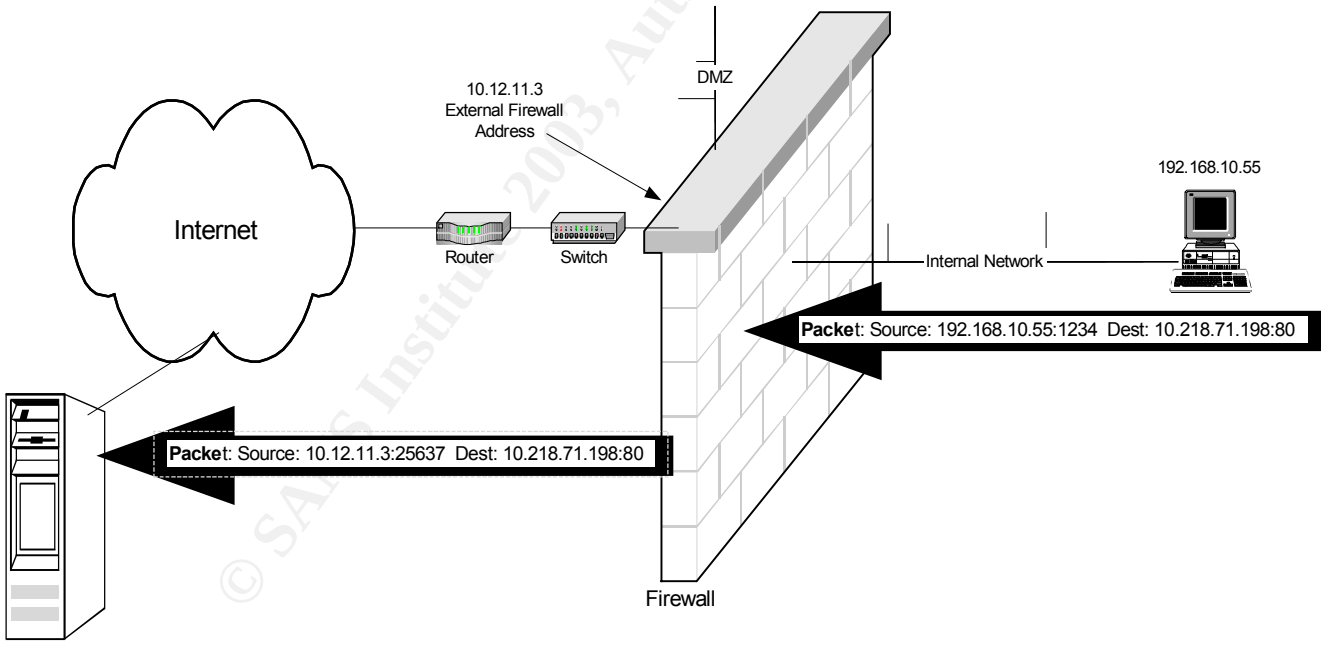

Address Translation Table

Internal Host Address	External Host Address	Firewall Translated Address
192.168.10.55:3567	10.218.71.198:80	10.12.11.3:25637

When the remote server responds it sends the packet to the address and port that the firewall furnished. The firewall checks its translation table for a matching entry and forwards the packet to the requesting host. See step 2.

Figure 6: Address Translation Step 2

Step 2.

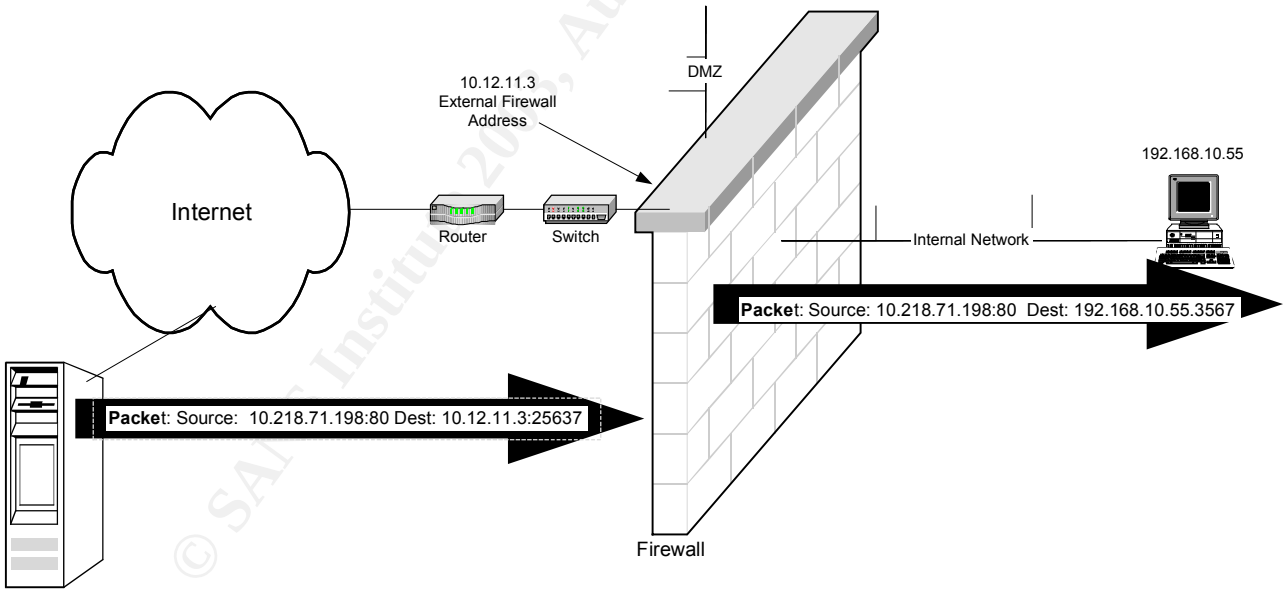

Address Translation Table

Internal Host Address	External Host Address	Firewall Translated Address
192.168.10.55:3567	10.218.71.198:80	10.12.11.3:25637

In this way NAT can mask the internal addresses for additional protection.

DMZ:

DMZ stands for De-Militarized Zone. The DMZ is an area between your inside network and the outside. It is an area that is less secure than the inside network but more secure than the outside.

Figure 7:DMZ one Firewall

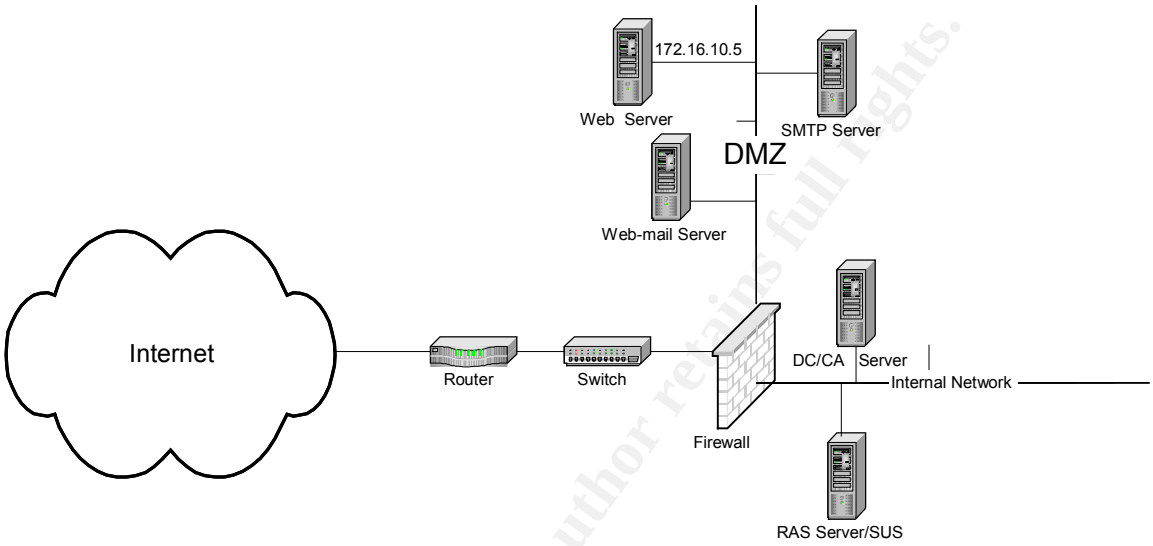

Figure 8 DMZ Two Firewalls

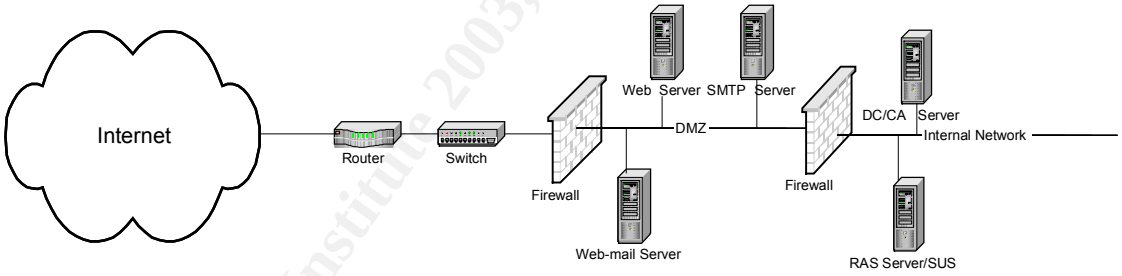

The DMZ allows you the ability to access systems like web servers and mail servers without allowing direct access into the inside network. For example a web server in the DMZ may be accessed by port 80 (HTTP) or Port 443 (HTTPS) from the Internet. That same server may be able to access a SQL server on the inside with its reserved IP address but there is no direct access from the outside to the SQL server. The reserved address cannot be routed on the Internet so it cannot be spoofed; the servers in the DMZ are especially hardened against attacks and are monitored closely.

Firewall Solution:

We have chosen a stateful firewall for our company from Cisco called PIX with 3 Ethernet interfaces one external, one DMZ, and one internal. That gives us the best protection without sacrificing performance.

The DMZ we created contains a web server with access to a application server for on-line sales, an smtp gateway to forward messages into the internal e-mail server, and a web server that allows users to access their e-mail via SSL.

Anti-Virus Software:

Anti-Virus Software is a way of protecting our network from viruses, worms, and malicious code that reach systems through the network or removable media.

Viruses:

Viruses come in three types program, boot record, and macro viruses. Program viruses target programs and get activated and spread when the program is executed. Boot infectors infect the boot record of a disk and are executed during the boot process where they are spread to memory and other disks. Macro viruses infect data files and are executed by programs like MS Word or Excel. They can be spread as e-mail attachments. Our sample environment is highly at risk to all these viruses, which is why we choose to implement an anti-virus policy. The method of protection we choose is two fold. First there is a file scanning based anti-virus protection scheme for our servers and workstations. Second there is an e-mail scanning virus protection scheme to protect against attacks though the mail.

What we are looking for is not only something that works but is easy to deploy and maintain. If we can't keep up with the anti virus definitions updates then we are still at considerable risk. Also we want to make sure that the product we have will be properly supported. There are several good tools out there including Symantec's Norton anti-virus, InoculateIT from CA, and NeatSuite from Trend Micro. All are products I've used and work well.

Anti-Virus Solution:

We choose Trend Micro's NeatSuite. NeatSuite gives us several nice tools to combat viruses at an attractive price. Your mileage may vary about these products but what I want you to take away from this is what the product does and how we use it. The first is file-scanning protection for our servers and workstations. Server Protect and Office Scan are programs that run on each server and workstation to scan incoming and existing files for viruses.

This package allows us to download the definitions to a central server and push them to our clients and servers. First we test the updates on our test server and some test clients. Then we allow the server to push out the updates automatically.

The ScanMail for Exchange checks all incoming emails for viruses heading to the exchange servers. It runs on our internal mail server and is updated automatically.

InterScan Virus wall is the tool we use on our smtp gateway. We scan incoming mail in the DMZ and on the internal server. That gives us two checks on incoming mail. We scan files as they are read on the workstations and all e-mail attachments on the exchange server. That gives us two checks on outgoing mail.

All in all we have an automated system that helps us to comply with our anti-virus Policy.

This picture gives you an idea of the type of anti-virus protection we require. The outer circle represents the fact that all systems are subject to file-scanning anti-virus protection. The gray colored areas depict the e-mail scanning anti-virus protected areas.

Figure 9: Antivirus protection

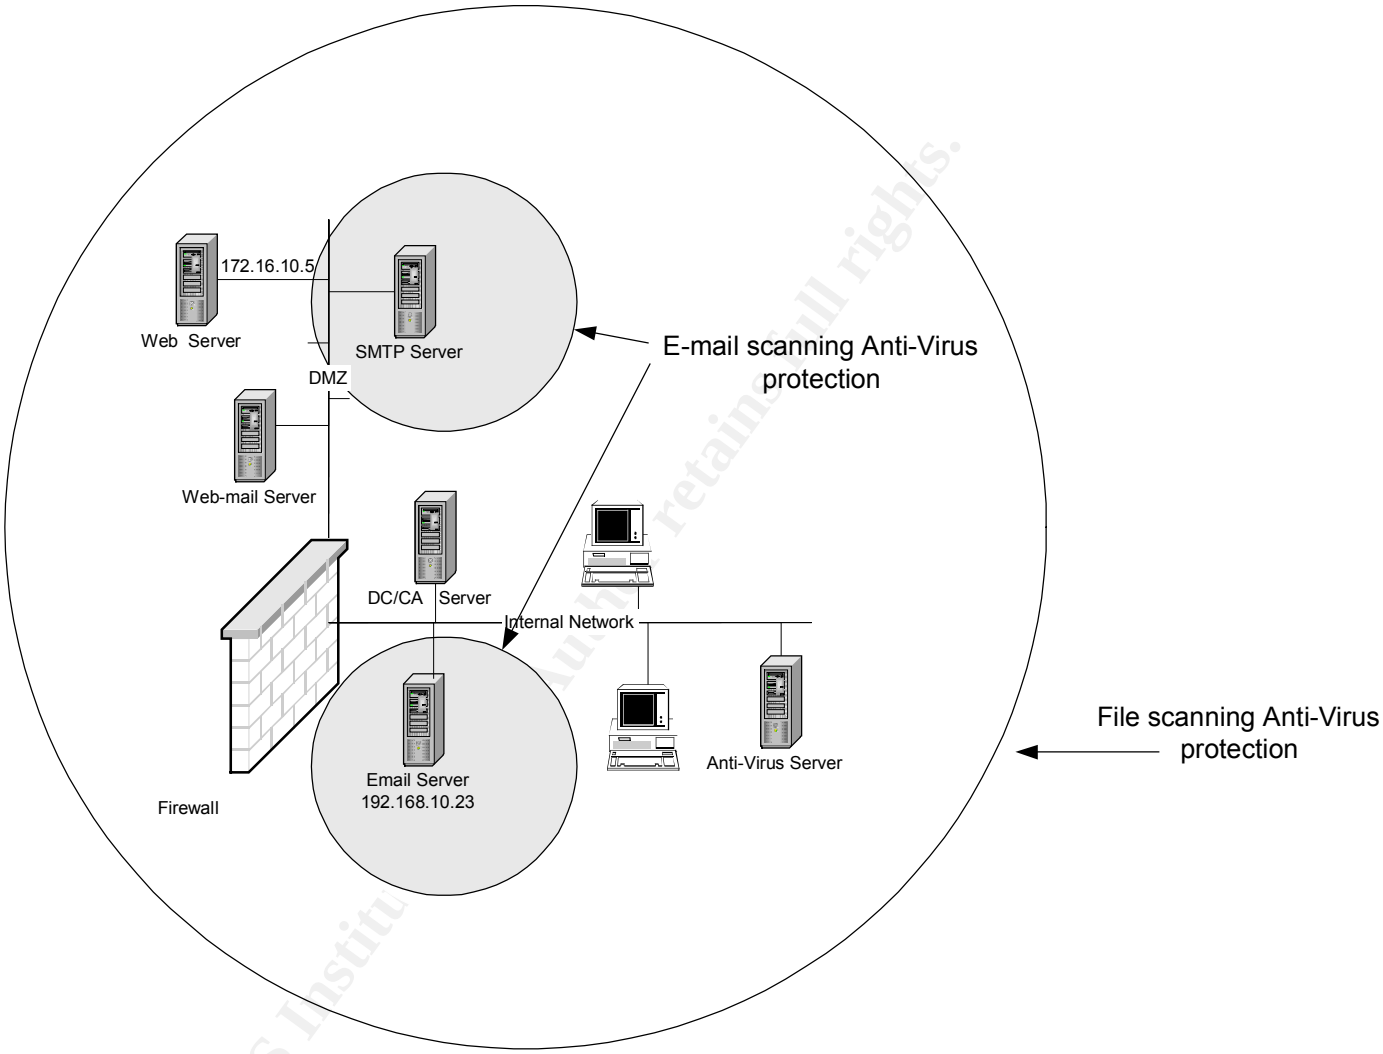

Authentication and Authorization:

Authentication and Authorization helps prevent unwanted virtual access to network assets.

Authentication is the process of validating that someone or something is who or what they say that they are. Authorization is the granting of access, rights or privileges based on a valid authentication.

Authentication solution:

In our network we use two-factor authentication. Two-factor authentication is when you combine something you have with something you know for proof of authenticity. We use SecureID from RSA and Microsoft's Active Directory.

With SecureID all of our remote access and VPN users are issued a token card. This card displays a code that changes every 60 seconds. This code is combined with a user defined PIN to complete the first part of authentication. Once the user has passed the SecureID portion they are then prompted for a Microsoft Active Directory username and password.

Authorization solution:

Microsoft's Active Directory also provides the authorization for network access. NTFS file permissions, combine with group policies, Access Control Lists (ACL), user rights, and share permissions to create the authorization mechanism. Active Directory manages and enforces all of these.

Password Protection:

Passwords are the keys to access in our network so they must be protected. Microsoft's Active Directory allows us to define password policy. Some examples of that policy are.

- Maximum Password Age – Length of time a password can be used before it must be changed.
- Minimum Password Length – defines the smallest number of characters a password can contain.
- Password History - requires a number of passwords that must be used before reusing the password. This keeps users from using the same 2 or 3 passwords over and over.
- Minimum Password Age – This is an amount of time that must expire before the use can change their password. This prevents the user from changing the password several times in succession to defeat the Maximum Password Age limit to use a favorite password.
- Password Complexity – requirements that the password must meet, like the use of alpha numeric and special characters.

We couple the password policy with an account lockout policy. The account lockout policy locks out users after 3 failed attempts to login for 15 minutes.

Patches, Updates, and Hotfixes:

By policy all of our systems must be as up-to-date as reasonably possible to prevent hackers form exploiting published vulnerabilities. Like most companies today, Microsoft publishes alerts about discovered weaknesses in their products' security. People who would exploit these weaknesses also read these alerts. So it is important to keep up with the updates.

Patches, Updates, and Hotfixes solution:

For this requirement we choose to use Microsoft's Systems Update Server (SUS). SUS is a free program located on the Microsoft Website.

http://www.microsoft.com/windows2000/windowsupdate/sus/default.asp

SUS includes a client and a server. The server component is installed on a computer running Windows 2000 Server. It connects to Microsoft update servers to download the latest updates. When the updates are downloaded, we test the updates in our environment and then decide which updates to approve for installation. These updates will not be made available to the clients until the security department approves them. The SUS server logs the update status of clients to our internal Intranet web server. That will allow us to track the status of updates. The server can be scheduled to deliver updates at times that are best for our users. This process allows us to make sure all workstations are being kept up to date.

The client portion is the automatic update service, which comes with Windows XP and Windows 2000 SP2. It will be triggered to update through group policy. This means that The local user will not have to be given local admin rights to be able to install the updates. By not giving the users that right we can more effectively lockdown the desktops thereby preventing user mischief and accidental destruction.

Security by Obscurity:

Security by Obscurity is a method of securing things by concealment. The idea is that if no one knows that you have a bundle of cash hidden in your house it is not likely to be stolen. Security by Obscurity has lots of holes. Secrecy is only effective against certain types of threats. Obviously this theory is of little use in case of a fire. Just because you believe it is secret doesn't mean that it is. You are probably not the first person to hide money under the mattress. This form of security is only good until someone uncovers the secret. So as a method of security it is weak.

Having said that remember that we are layering security methods as a part of a defense in depth strategy. Treating information about your network, servers, applications, and processes on a need to know basis is a sound practice.

According to the confidentiality policy of Goodman Inc. configuration information like network diagrams, server configurations, software revisions, patch levels, and security policies themselves are considered confidential materials with rules governing their creation, access, and disposal.

With this policy we are not trying to stop all of the super hackers of the world. We just want to place a roadblock in the way to stop some of the attempts.

Security by Obscurity solution:
We created a confidentiality policy to limit access to documentation on a need to know basis.

Virtual Private Network:

Virtual Private Networks or VPNs are encrypted pipelines passing though outside networks like the Internet to connect remote systems and networks to the inside network. Encryption is a process of encoding information so that it is only readable to those who can decrypt it. One type of encryption is called private key encryption. Private key encryption is the process of encoding a message with a key that can only be unlocked with that key. This means that both the encoder and decoder must have the same key. Another type of encryption called Public Key Private Key requires two keys. The message is encoded with one of the keys but it can only be decoded with the other key. The public key is made publicly available. The encoder of the message holds the private key.

There are a lot of different configurations of VPNs; the two we will be concerned with are IP LAN to LAN and IP host to LAN.

LAN-to-LAN VPN

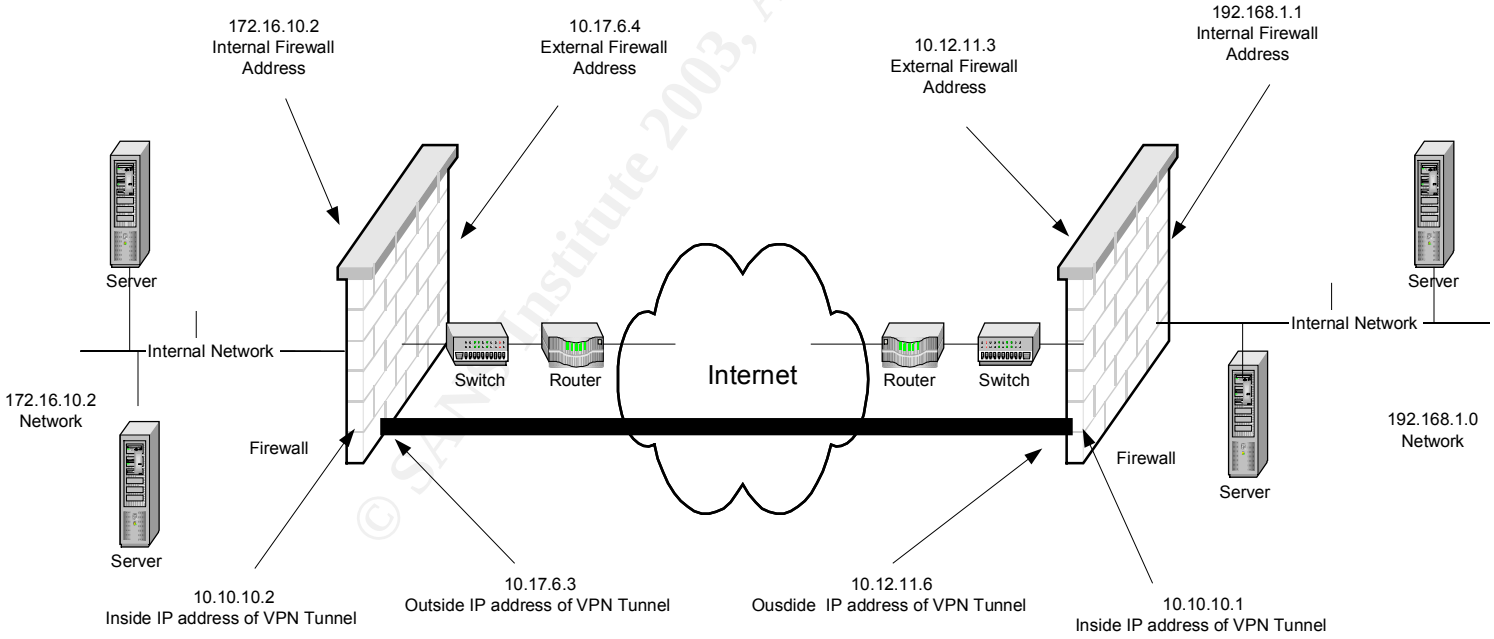

LAN to LAN VPN:

With LAN to LAN you have two LANs connected via an encrypted tunnel. Two devices, one at each end creates the tunnel. These devices are usually routers, systems, firewalls, or special VPN appliances. The tunnel is established between the devices using a tunneling protocol. The two most common tunneling protocols are PPTP (Port to Port Tunneling Protocol) and L2TP (Layer Two Tunneling Protocol).

In the diagram you have two LANs separated by the Internet. Access between the two networks is provided via the 10.10.10.0 network. The external IP addresses of the tunnel are transparent to systems on the LANs. Traffic heading from 172.16.10.2 to 192.168.1.45 would pass from the default gateway to the opening of the tunnel 10.10.10.2. It would be encrypted and passed to the external interface 10.17.6.3, which passes it to the Internet router. It then gets handled like any other Internet traffic and ends up at 10.12.11.6. The VPN device decrypts the packet and passes it to the gateway for the 192.168.1.0 network that routes it to 192.168.1.45. The entire VPN operation is transparent to the end devices.

Host to LAN VPN:

Another type of VPN configuration consists of individual devices each having their own VPN to the LAN. These are typically to users who work from home or small offices that have connections to the Internet via DSL or Cable modem. Technically these work much in the same fashion but there are some considerations. LAN-to-LAN VPNs typically have firewalls to protect them from the Internet. It is just as important to have firewalls on the individual systems as it is for the LANs.

Without a firewall on your system you run the risk of opening your entire LAN to wide open access from the Internet. The reason is simple; once you connect that VPN to your LAN anyone that has compromised that PCs security now has access to your LAN through the VPN.

It is for this reason that we established a strict VPN policy. It includes the requirement that anyone that uses a host to LAN VPN must sign an agreement to follow this policy.

PPTP:

PPTP is encapsulated and encrypted PPP on IP. PPP is the Port-to-Port Protocol and was originally developed to encapsulate data and carry it over point-to-point links. PPTP is basically a vehicle for carrying PPP. The PPTP forum consisting of Microsoft Corporation, Ascend Communications, 3COM, ECI Telematics, and US Robotics created PPTP.

PPTP operates in a client/server fashion. The server is called a PPTP Network Server and the client is a PPTP Access Concentrator. A client/server pair creates a tunnel. PPTP relies on TCP/IP and uses an enhanced GRE (Generic Routing Encapsulation) mechanism for carrying PPP packets between the server and client. (RFC-1701). GRE is a transport layer encapsulation protocol. It authenticates the source of a packet by the use of a key field in the GRE header that indicates which session a particular PPP packet belongs to. This allows for multiple sessions through the same tunnel.

PPTP Packet

Data-Link Header	IP Header	GRE Header	Encrypted PPP Payload	Data-Link Header

The authentication protocol in Microsoft PPTP is MS-CHAPv2. It uses a hash of a random number challenge, a peer response challenge and the client's username to authenticate.
The encryption protocol for Microsoft's PPTP is Microsoft Point to Point Encryption (MPPE). MPPE uses private key encryption with a separate key for each direction of the conversation. MPPE keys can be either 40 bits or 128 bits and are created using a hash of the user's password. All passwords can be cracked given enough effort and weak passwords are easy to crack. Because of this PPTP is not considered to be as secure as IPSec over L2TP.

L2TP:

L2TP is a combination of PPTP and Cisco's L2F tunneling protocol.
L2TP is a layer 2 transport protocol and must be accompanied by some form of encryption usually IPSec. L2TP can use IPSec to encrypt PPP frames.

IPSEC:

IPSec is a layer-3 encryption protocol. By functioning at the network layer IPSec is able to encrypt TCP and UDP packets. IPSec is made up of three protocols. Authentication Header (AH), Encapsulating Security Payload (ESP), and the Internet Security Association Key Management Protocol (ISAKMP).

ISAKMP:

The Internet Key Exchange protocol (IKE) uses ISAKMP to Authenticate IPSec peers, create keys for encryption and establish Security Associations. (SA) IKE is a protocol which implements the Oakley key exchange inside the ISAKMP framework.
IKE uses a pre-shared secret key or public keys to authenticate and establish a session key. After the exchange, the remainder of the session is encrypted with the session key. This session key can be changed or refreshed periodically to protect against it's being compromised.
A Security Association is the policy that defines the connection. It is established and agreed upon at the creation of the connection. It contains parameters like encryption algorithms and encapsulation types. For example on our PIX firewall we set the parameters for 3DES-168 = Triple-DES encryption with a 168-bit key. If one of the parties trying to connect doesn't agree on the encryption then no SA will be established.

AH:

Authentication Header is a way to validate who the entity is that sent the packet. AH is applied to an outbound packet only after an IPSec implementation determines that the packet is associated with an SA that calls for AH processing. The AH uses the session key and an authentication algorithm to create an Integrity Check Value (ICV). Using

almost the entire packet excluding some fields that change like the TTL field of the IP header. It stores this ICV in the AH which is inserted into packet between the IP header and the data. When the packet arrives at the end of the tunnel this action is performed again. The ICV created at the end is compared to the ICV that was sent with the message. If they are different the packet is rejected.
The data in the packet is not encrypted. This process is used to authenticate the source of the packet and that it has not been modified.

ESP:

Encapsulating Security Payload is where the meat of the encryption takes place. After the SA is established IPSec uses the session key to encrypt the data of the packet. The packet traverses the tunnel is decrypted and passed to its destination. ESP has two modes of operation, transport mode and tunnel mode.

In transport mode an ESP header is inserted between the IP and TCP headers of an IP packet. The TCP header and payload are encrypted according to the SA. An ESP trailer and an ESP authentication field are added to the end. The trailer contains padding bits to extend the packet to a fixed length. While the ESP authentication contains a hashed value to authenticate the packet. Transport mode is typically used as an end-to-end solution, where the encryption and decryption take place at the source and destination respectively.

Tunnel mode encapsulates the original IP packet inside of another IP packet. The entire packet is then encrypted. Tunnel mode is typically used in a LAN-to-LAN configuration, where an encrypted tunnel connects two networks.

Solution:

LAN-to-LAN Solution

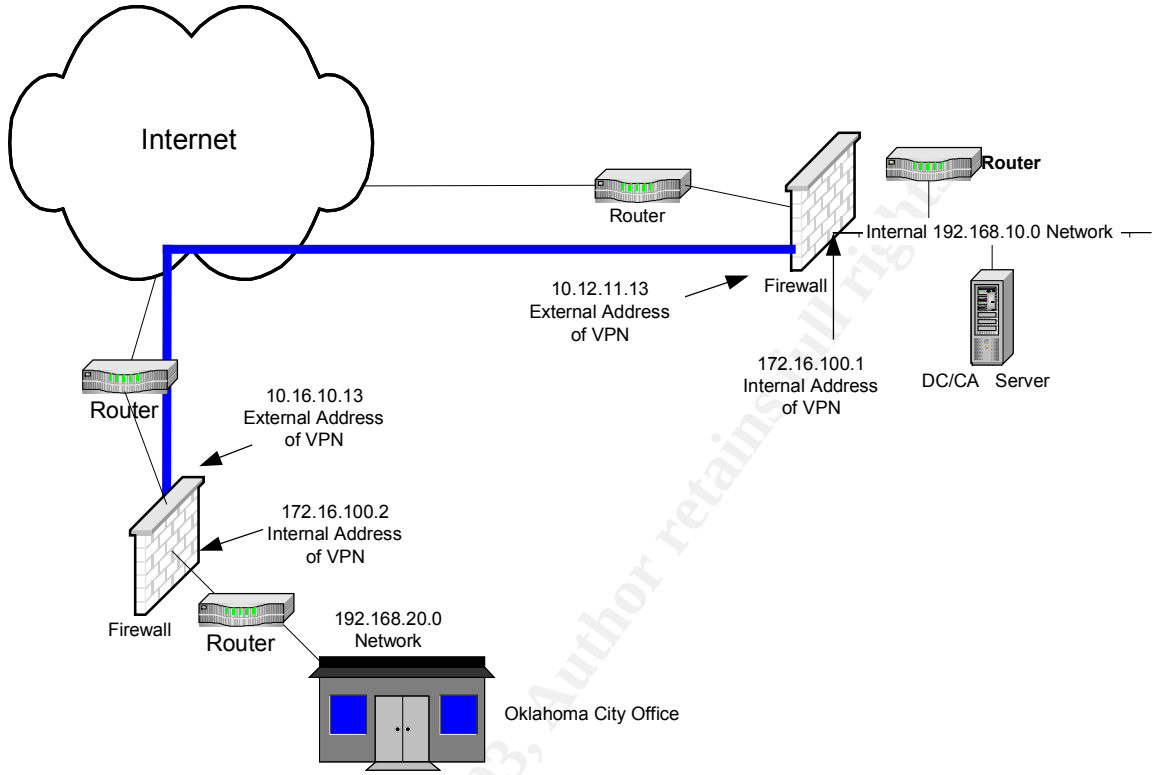

Since we use PIX firewalls at both of our sites we created a LAN-to-LAN tunnel from one firewall to the other using L2TP and IPSEC. The tunnel consists of two external IP addresses and two internal Private IP addresses. We made entries in our internal routers forwarding traffic destined for the remote network to the internal address of the tunnel.

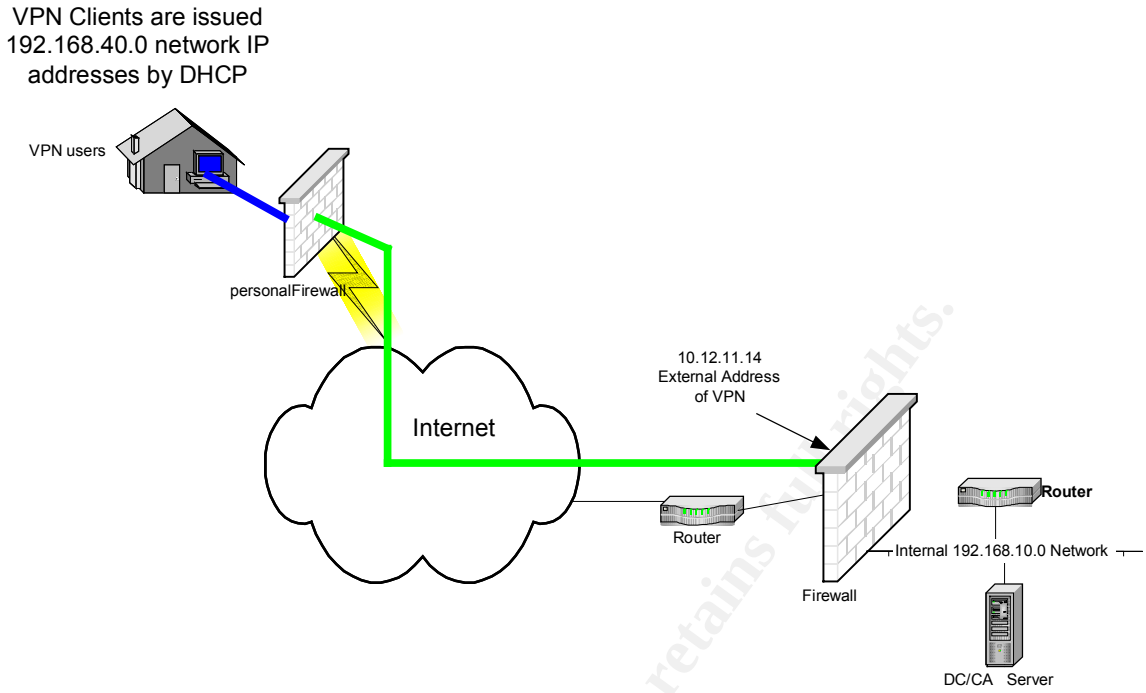

For the Home Office users we established a Host-to-LAN VPN using the PIX firewall and the Cisco VPN client. The client is loaded onto the remote PC and includes a firewall that is automatically engaged each time the VPN is established. The VPN clients are issued Private internal IP addresses via DHCP. They are allowed access to Internet sites only by using the Goodman Inc. network for as long as the VPN is connected. Once the VPN connection is dropped Internet activity resumes through the user's ISP.

Remote Access Service:

For those users that travel or do not have a home Internet connection we provide Remote Access Services (RAS). RAS services are only allowed to laptop users by policy. These laptops must all adhere to the Remote/Dial in security policy. This policy contains rules governing the use and security of dial-in capable systems.
The RAS server must comply with the Authentication policy.

Solution:

We use a Microsoft Windows 2000 server running RRAS with Equinox Multi-Modem adapter cards. The server requires SecureID authentication prior to Windows Active Directory authentication.

Physical Security:

Let's not forget the obvious. If your machine is easily accessible then it is vulnerable to all sorts of attacks. Obviously we want to keep our servers and network devices in a locked temperature controlled room. We want to limit access to only those people who

need access directly to the servers and network devices. That means that we don't want printers, office supplies, or other such things located in the locked room. We use a coded combination protected door. We don't want the cleaning or maintenance people having free run in there. If anyone other than authorized IT staff must enter the room an authorized IT staff member must be present.

Physical Security solution:

Rules about physical security are in several of our security policies. For example our Network Switch and Router Security Policy as well as our Server policy requires that these devices be in a locked computer room or communications closet with IT authorized access only.

Security Awareness:

Of the tools available for securing your network one of the most overlooked is security awareness of your users. If users understand the possible threats involved and the harm they can cause, they can help you to defeat them. For example everyone should know not to open e-mail attachments from people they don't know. But do they know what to do if they get one?

To combat this we created a program to make everyone more aware of security issues. Here are some of our objectives and solutions.

- Security Orientation – All employees are required to watch a short video made by the security department.
- Lunch and Learn – Twice per year the company holds a security awareness meeting, lunch is provided.
- Highly visible contact list – We hand out sticky-note pads with the security incident response team's contact info on it.
- A Strong Message – We stress that security is everyone's responsibility.
- Discuss the penalty of failure – Security awareness training includes estimated costs caused by failure of security measures.
- Logon Warning Messages – Messages alerting people that this is a private network and that unauthorized usage will be penalized.

The idea of the program is to get people thinking about security and knowing what to do in case of an event.

How do I know if the protection is defeated?

Our third question, "How do I know if the protection is defeated?" is a good one. Let's say you were to hire someone to watch your house while you were out of the country. Then you called them to check on things. Would it be acceptable for them to tell you "Well the day you left I locked all the doors and I haven't heard of any problems"?

It's doubtful that would be the answer you were looking for. You would probably want to know more things like. Is the house still there? Are the things all still inside it? Is it damaged? That's the point of hiring someone to watch it.

So who is watching your network? Now that we've locked the doors how do we tell if someone broke in through a window? The Answer is auditing, logging, and intrusion detection.

Auditing:

Auditing is the process of tracking certain types of events. Windows 2000 domains with Active Directory can perform auditing. This auditing is not enabled by default and must be enabled. You set the auditing levels based upon the requirements in the Audit policy you created in the earlier steps. There is a step-by-step document for enabling Windows 2000 systems for auditing from Microsoft titled "HOW TO: Enable and Apply Security Auditing in Windows 2000". The URL is http://support.microsoft.com/default.aspx?scid=kb;en-us;300549#2.

You enable auditing locally on each system or network wide by use of group policy. Notice in example 10 under Local policies there are three folders one of which is Audit Policy. In the right window there are events that you can audit.

Figure 10: Audit Policy

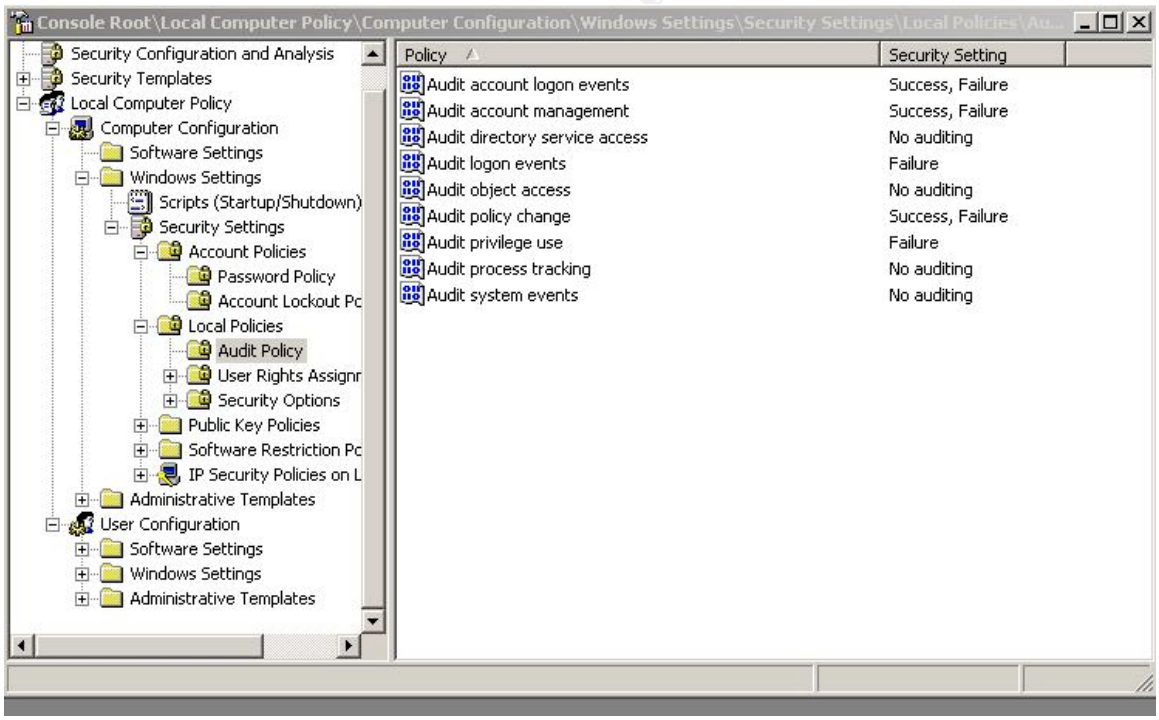

Figure 11: Auditing

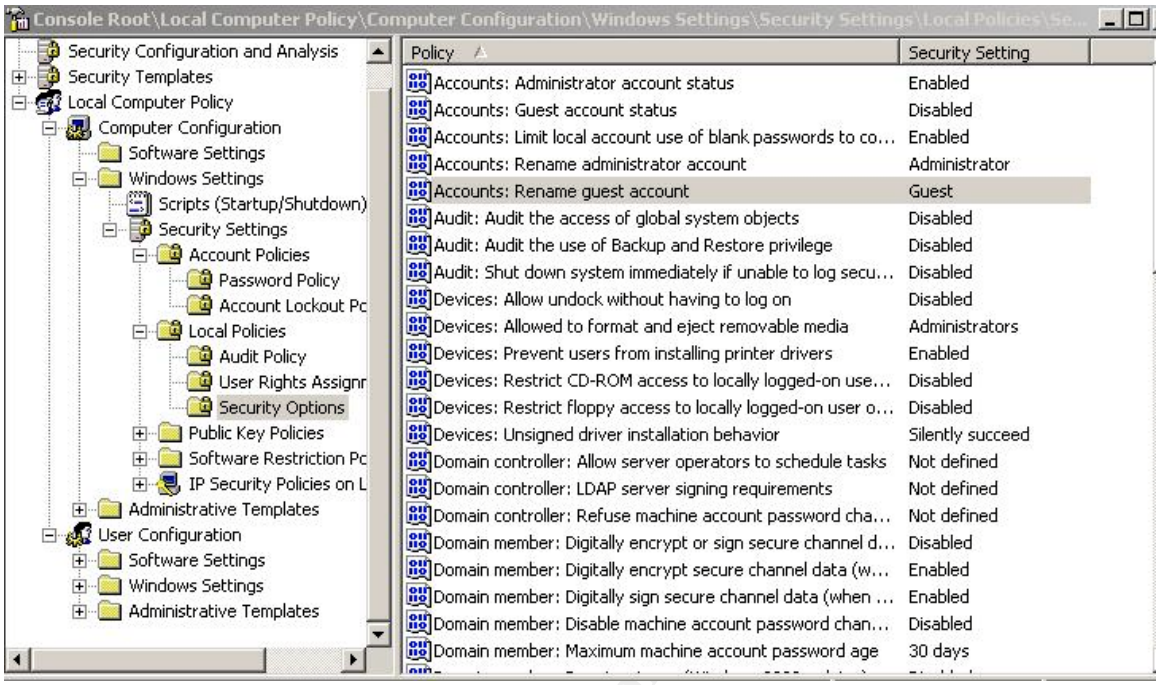

Notice in figure 11 there are three settings that can be enabled for auditing. The third one will if enabled shutdown the system if it cannot log security events. This may seem drastic but could be necessary depending on the level of security required for the system. You can also set Cisco devices to track certain types of events. By default most Cisco devices are set to log major events (severity level –3 errors) to the console. The level of tracking as well as the destination of the logs is configurable. On most Cisco devices switches, routers, and firewalls there are facilities like TCP, IP, TFTP, Telnet etc that can be monitored each with eight severity levels to which logging can be enabled. 0 through 7 with 0 = emergencies, 1 = alerts, 2 = critical, 3 = errors, 4= warnings, 5 = notifications, 6 = informational, and 7 = debugging. Through commands on the devices you enable the level of messages you wish to watch.

Example:

Console> (enable) **set logging level all 5**

All system logging facilities for this session set to severity 5(notifications)

Console> (enable)

From "Cisco 7600 Series Routers Configuring System Message Logging"
http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a008007dbc9.html - xtocid50701

Auditing works together with logging to provide critical information about the system and Network.

Logging:

Logging is the actual writing of events to some type of storage for retrieval at a later time. Microsoft uses the Event log to save most pertinent information. There are three main logs that come by default with Windows 2000. They are application, security, and system. There are other logs that are installed depending on which applications are installed as well. The security log is where audit information will be held.

Figure 12:The Security Log

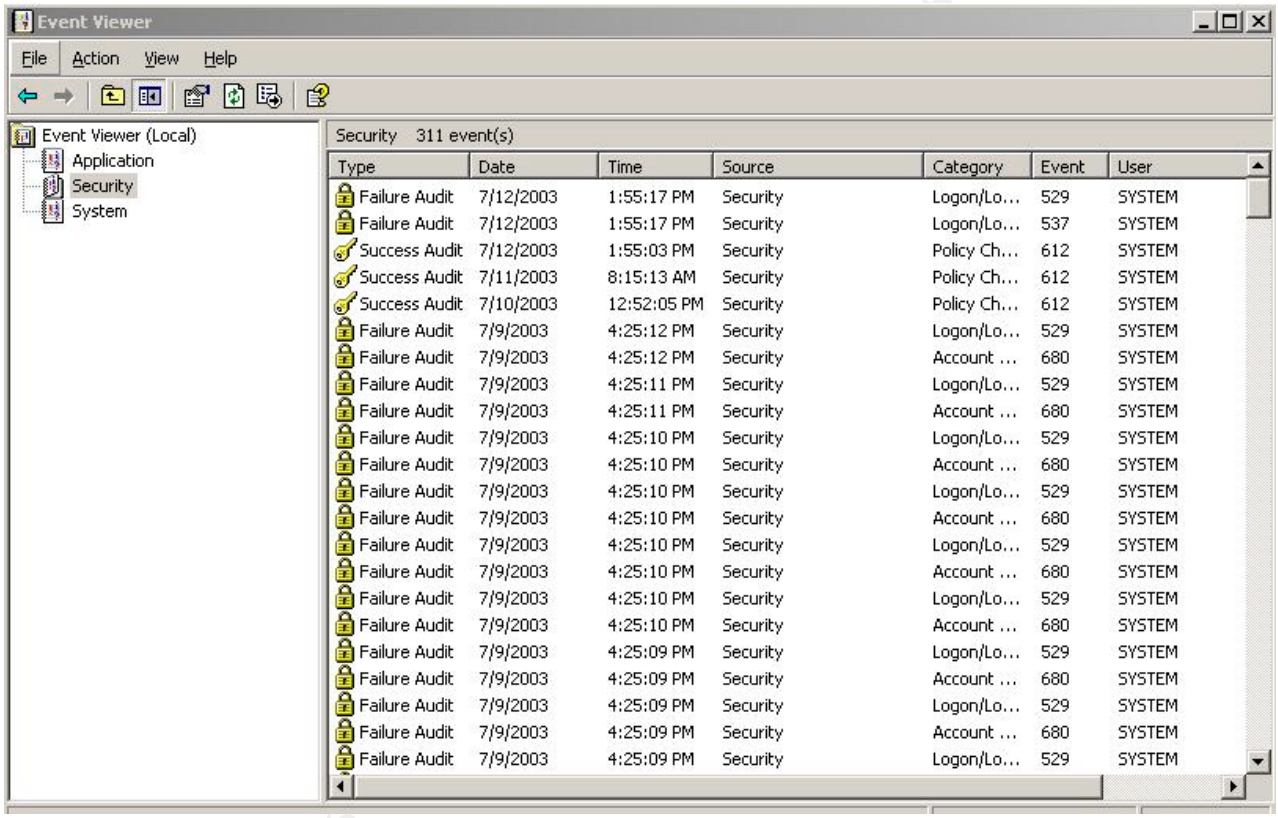

Type	Date	Time	Source	Category	Event	User
Failure Audit	7/12/2003	1:55:17 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/12/2003	1:55:17 PM	Security	Logon/Lo...	537	SYSTEM
Success Audit	7/12/2003	1:55:03 PM	Security	Policy Ch...	612	SYSTEM
Success Audit	7/11/2003	8:15:13 AM	Security	Policy Ch...	612	SYSTEM
Success Audit	7/10/2003	12:52:05 PM	Security	Policy Ch...	612	SYSTEM
Failure Audit	7/9/2003	4:25:12 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:12 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:11 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:11 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:10 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:09 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:09 PM	Security	Logon/Lo...	529	SYSTEM
Failure Audit	7/9/2003	4:25:09 PM	Security	Account ...	680	SYSTEM
Failure Audit	7/9/2003	4:25:09 PM	Security	Logon/Lo...	529	SYSTEM

Figure 12 is an example of the Windows Security Log showing success audits with a key symbol and Failure audits with a lock. Figure 13 show an example of one of the audit messages.

Figure 13: Example of an Audit Message

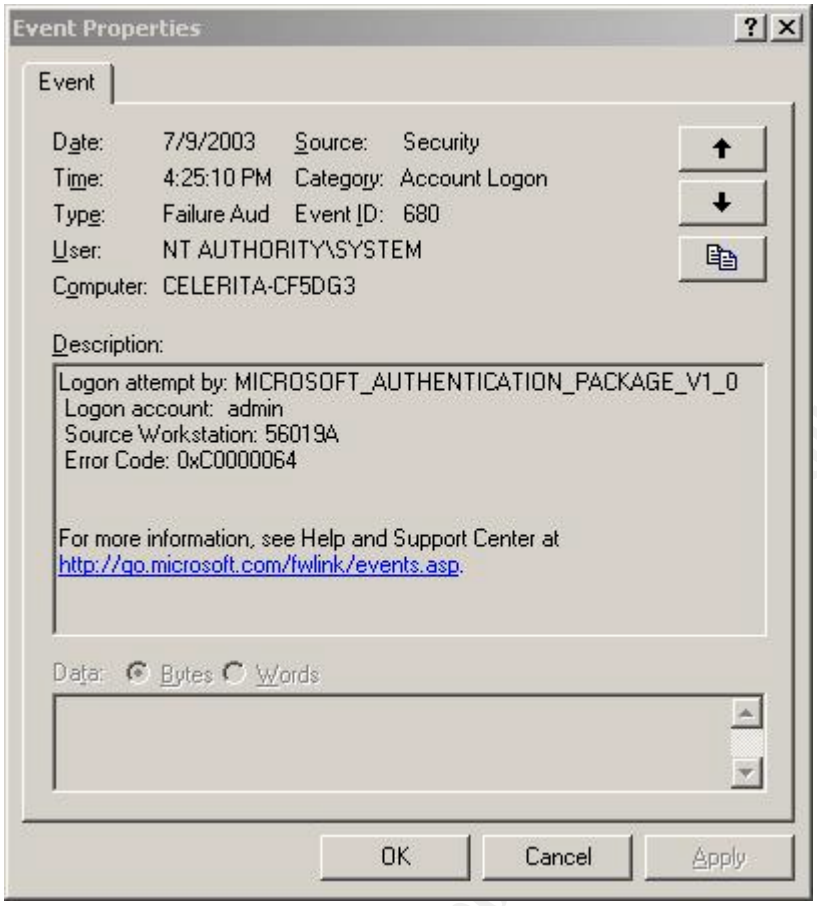

This example shows a failed attempt to logon to the admin account from workstation 56019A. This type of information can be useful in determining who is trying to access our systems.

Cisco devices can also log information locally or to remote systems. You can direct these devices to log to a Unix based syslog server. The following commands are from a Cisco router configuration but are similar across many Cisco devices.

Example:
Console> (enable) **set logging server 10.10.10.100**
10.10.10.100 added to System logging server table.

Console> (enable) **set logging server facility local5**
System logging server facility set to <local5>

Console> (enable) **set logging server severity 5**
System logging server severity set to <5>

Console> (enable) **set logging server enable**
System logging messages will be sent to the configured syslog servers.

Console> (enable)

From "Cisco 7600 Series Routers Configuring System Message Logging"
http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a008007dbc9.html - xtocid50701

Auditing and Logging Solution:

Our company implements Windows system auditing with Microsoft's Active Directory. We audit for specific events as follows.

Auditing and Account Policies Requirements

Audit Policy (minimums)

All passwords are **at least 8 characters long** (minimum).
All passwords are **no more than 90 days old** (maximum).
Audit Account Logon Events: **Success and Failure**
Audit Account Management: **Success and Failure**
Audit Directory Service Access: **Not Defined**
Audit Logon Events: **Success and Failure**
Audit Object Access: **Failure (minimum)**
Audit Policy Change: **Failure (minimum)**
Audit Privilege Use: **Failure (minimum)**
Audit Process Tracking: **Optional**
Audit System Events: **Success and Failure**

Account Policy

Minimum Password Age: **1 day**

Maximum Password Age: **90 Days**
Minimum Password Length: **8 Characters**
Password Complexity: **Enabled**
Password History: **24**
Store Passwords using Reversible Encryption: **Disabled**
Account Lockout Policy
Account Lockout Duration: **15 Minutes** (minimum)
Account Lockout Threshold: **3 Bad Login Attempts** (maximum)
Reset Account Lockout After: **15 Minutes** (minimum)

These events are logged to the event logger. The Microsoft Windows Event logger handles initial logging. Adiscon's WinSyslog and EventReporter augment this process. The event logger logs the events triggered by system events and the audit policy. Each server keeps it's own event log. One of the things a hacker is prone to do when they get access is to clear event log entries that may belie their presence.
WinSyslog is a Windows version of a Unix syslog server. EventReporter is a program that forwards Windows Events to a syslog server. Used together we have a created centralized logging system that makes it harder for hackers to cover their tracks. Each Windows server forwards events to a WinSyslog server; this server runs another Adiscon product called Monilog. Monilog creates daily filterable reports that can be published to a web site or emailed to administrators.

Intrusion Detection:

Intrusion detection consists of various ways to check for the existence of anomalies that often occur due to the presence of hackers. Port scanning would be an example.
There are two types of intrusion detection host based and network based.

Host based IDS:

Host based intrusion detection systems (IDS) monitor their network connections and file system status. The logging process we just talked about is a form of intrusion detection.
Programs that monitor for file system integrity are another form of Host based IDS. An example is Tripwire for Windows. Tripwire will create a cryptographic hash for the files that you want to monitor. It compares the stored hash with the hash created when the file is checked. When the files are modified it will notify you of the changes.
A free tool that can be used for a similar purpose is WinDiff. The WinDiff tool takes two files or two directories and does a byte-for-byte file comparison. It then displays the results in a scrollable window, lining up the identical parts, and marking with color the parts of the files that differ. You can also select to view the differences as in the example in Figure 14.

Figure 14:WinDiff

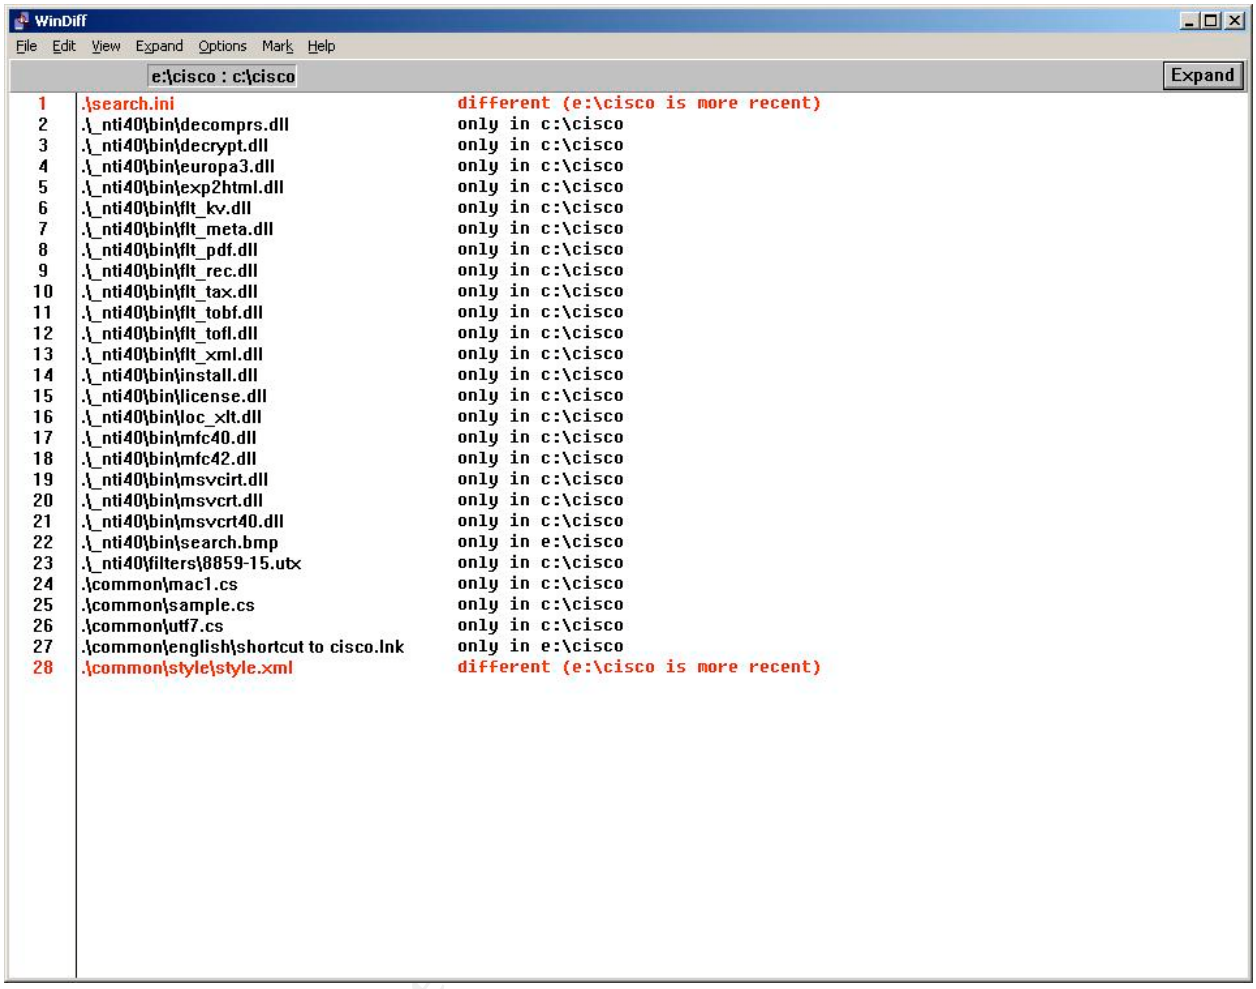

This example shows only the files that are different between the two directories and how they differ. This provides a way to compare files and directories to known entities to see if they have been changed. Windiff.exe is part of the Windows 2000 Resource Kit.

Network Based IDS:

Network Based IDS usually consists of a system that scans network packets audits packet information, and logs any suspicious packets. They sit in a location that allows them to monitor all traffic coming in form the Internet. This typically requires their network interface to operate in promiscuous mode.
Snort is network based intrusion detection system, capable of performing traffic analysis and packet logging on IP networks. http://www.snort.org/

Companies like Oculan make appliances that sit on your network and provide monitoring and intrusion detection. The Oculan 250 is an Intel based system that runs Linux and Apache web server. The system can monitor 25 devices (servers and network devices) and 250 desktop systems. The device is Plug n Play, and with a web interface for management and is fairly easy to setup.

IDS Solution:

For a small network like the one Goodman Inc has, money is always an issue. Though a host-based solution probably offers the best protection it is also labor intensive. It requires that every system be checked and their file systems must be compared to known good ones that must be updated each time the system changes.

For these reasons we chose to implement a centralized logging system using Adiscon's Event Reporter and WinSyslog. We also chose to implement the Oculan 250. At around ten thousand dollars it is a total IDS solution for about the same cost as a single application server.

This coverage gives us notifications of suspicious events and a single source to review system log information.

What do I do if the protection is defeated?

Security Incident Response Team:

A critical piece of any security design is having people that know what to do in case of a security breach. A security response team can act as a single point of contact for security issues as well as develop, maintain, and track incident handling and reporting. The team should understand the issues involving security. They should be the ones that create the incident response plan. They should also be the focal point in driving post-incident activities. For example after the discovery of malicious code on a system the security response team may request that a particular system be removed from the network until it is restored from backup to a pre-incident state by the IS department.

There is a good article on creating an incident response team at the Cert. org website.
http://www.cert.org/csirts/Creating-A-CSIRT.html

In a small organization the team may only be one or two members. The security training may also be limited. That is why it is important to have an incident-handling plan.

Incident Handling:

The incident response team should develop a plan to handle security incidents. This plan should include guidelines concerning who should be contacted and when during the aftermath of the incident. It should list internal and external contacts of vendors, consultants, technical resources, and anyone that may need to be contacted to resolve the situation.

The plan should include a series of steps to follow for specific types of issues. For example if you are dealing with an e-mail virus the first step might be to remove the e-mail server from the network until you can isolate and deal with the virus.

The plan should also include an escalation process. The last thing you want is to notify people at the top of your organization with bad news but waiting too long can be disastrous. For example if you are going to take 4 hours on Saturday morning to rebuild a system that was just whacked. It might be important to notify the 200 users that were planning on working overtime on that system that they need not report to work. The management may not like the fact that the system was down, but they really don't want to pay 200 workers overtime just because they weren't told not to show up.

The plan should also layout post-incident steps that need to be taken after recovering from the incident. These should include a review of what failed, a review of how the incident was handled, and what should be done to prevent similar events in the future. This process could be as simple as filling out a Security incident report form.

Restoring systems:

Part of your security incident handling should include your disaster recovery plan. Depending on the incident it may be necessary to recover entire systems or data centers. So it is important to have a good recovery plan to protect your company's assets. Michigan State University has some good information on how to handle Disaster Recovery Planning of which can be found at the following website.
http://drp.msu.edu/index.htm
Including a Basic DRP outline of which can be found at the website.
http://drp.msu.edu/Documentation/1_rww-stepsmemo-feb2001.doc .

This subject cannot be covered in this paper I only wish to note its importance to a successful incident handling process.

Another important concern with Backup and Recovery is to be aware of the security implications of the backup media themselves. Are they securely located offsite? Do know what type of environment they are stored in? Is it protected from electromagnetic interference or could a local magnetic field from a nearby source damage your backups? What about the handling of them from your site to the storage?

One way to find out is to check for your self. Visit the storage site from time to time. Take note of what things are in the vicinity. Generators, motors or anything that can create a magnetic field are a problem. Besides that check your backups from time to time with a practice restoration of a system. This not only tests your backup media, but your backup and restore process as well. It is also a good way to make sure that your people understand what to do so there is no mistake at crunch time.

Learning from Failure:

Failure in the world of computer security is not only possible it is a virtual certainty. There will be holes you didn't expect. There will be viruses that your anti-virus doesn't detect. There will be threats that you cannot 100% protect your self from. That is the idea behind defense in depth when one method fails another must be defeated.

Another of these layers is your knowledge and ability to learn from your failures. The first time you picked up a hot pan and burned your hand you learned not to do that again without using a towel or pad.

That is why it is important to document security incidents, review their resolution, and determine ways to prevent them.

You are not limited to your own failures there are many places that publish security alerts. Theses alerts are usually based on the failure of someone else's security and are sent as a warning. There are several websites that routinely alert you to these types of failures to help you protect yourself.

Here are a few.

- http://www.cert.org/
- http://www.sans.org/newsletters/
- http://www.nipc.gov/

Your security team should routinely these or sites like them to keep up to date.

Document, Document, Document:

As hard as it is to keep your site protected from threats without proper documentation its almost impossible. Network diagrams, system information, disaster recovery information, incident reports, policies, and contact lists are all important documents to maintain. With an accurate network diagram I can better tell where I am susceptible to a specific type of threat. With accurate system documentation, I know which systems the slammer worm might damage. I can use incident reports to help me decide if I need that new intrusion detection system. My security policies create the ability to punish violators and my contact lists save my bacon when its in the fire.

It is not hard to keep up documentation it just requires discipline. The question you should answer is. Can you afford not to have it? Make sure that the people who need it have access to it but treat it on a need to know basis. Remember security by obscurity is a layer of defense.

Solution for What to do if the protection is defeated:

We created an incident response team for handling breaches of security. Also all members of the team are responsible for creating and maintaining the incident response procedures.

We make everyone aware of the team and what to do in case of an incident. One way we get the word out is we hand out sticky-note pads with the security incident response team's contact info on it at the bottom. That way people see how to contact the team every time they use a sticky note.

We use computer incident report forms. These include information about date, system, type of incident, damage caused, resolution, and ways to prevent future similar incidents.

We created a disaster recovery plan. This plan can be used to recover from catastrophic events.

We review our incident reports to help create protection against future events. We also use several websites to keep us informed of new alerts.

We keep accurate documents for our systems and networks and review them on a regular basis.

Summary:

To secure our network we had to answer four main questions "What do I want to protect?" "How do I protect it?" How do I know if the protection is defeated?" and "What do I do if the protection is defeated?"

We started with an inventory and risk assessment to determine what to protect. Then we developed policies and guidelines to define the levels of protection. We want to create a layered defense in keeping with our Defense in depth philosophy.

Knowing what we need to protect and our philosophy we looked at methods of protection. These included a firewall and DMZ to protect our network from the outside. Similar to the walls around a castle with outer walls surrounding inner walls. We employed an anti-virus solution to act as our individual armor. We deployed a two-factor authentication tool to validate our users. This would be like having a password to enter the castle and special areas within the castle.

We use tools to maintain our patches and service packs. This could be analogous to keeping the castles soldiers up to date on the latest warfare techniques.

We have VPNs connecting our remote site and users to our network with encrypted tunnels. We also have a remote access server for dialup users. Both of these use the two-factor authentication for strengthened protection.

We physically secure our servers and network devices from unauthorized access.

We perform audits and checks to discover if our security has been breached. If it has we have a response team and documented processes for handling the situation. If the event requires completely rebuilding systems we can do it.

We don't look at these events as personal failures but instead use them to learn how to prevent similar attacks in the future. We try to have multiple defenses so that if one fails another takes its place.

Security in my mind is synonymous with "pain in the neck". It makes it more difficult to use your assets and slows down everything. But just like having a key to start your car or open the door to your house, it is a necessity. You wouldn't think of using curtains as the only doors on your house or a push button to start your car. By the same reasoning in today's world you cannot leave your network open to the world. Security is a necessary anti-evil.

Citations:

Meritt, James W. A Method for Quantitative Risk Analysis
Computer Security Resource Center (CSRC) The URL is:
http://csrc.nist.gov/nissc/1999/proceeding/papers/p28.pdf

Bass, Tim and Robichaux, Roger Defense-In-Depth Revisited: Qualitative Risk Analysis Methodology For Complex Network-Centric Operations.
The URL is
http://www.silkroad.com/papers/pdf/archives/defense-in-depth-revisited-original.pdf

"The SANS Security Policy Project"
The URL is
http://www.sans.org/resources/policies/

Microsoft Solution for Securing Windows 2000 Server
Chapter 6 - Hardening the Base Windows 2000 Server
Microsoft Corporation www.Microsoft.com The URL is
http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodtech/Windows/SecWin2k/06basewn.asp

Rekhter, Y
"RFC1918 - Network Working Group Request for Comments: 1918 "
The URL is
http://rfc.sunsite.dk/rfc/rfc1918.html

Egevang, K and Francis, P
"RFC 1631 - Request for Comments: 1631"
The IP Network Address Translator (NAT)
http://www.faqs.org/rfcs/rfc1631.html

Postel, J
"RFC 796" September 1981
Network Working Group Request for Comments: 796
Address Mappings
The URL is
http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc0796.html

"Software Update Services" Microsoft Corporation
http://www.microsoft.com/windows2000/windowsupdate/sus/default.asp

"HOW TO: Enable and Apply Security Auditing in Windows 2000"
Microsoft Knowledge Base Article – 300549 Microsoft Corporation
The URL is http://support.microsoft.com/default.aspx?scid=kb;en-us;300549#2.

Minasi, Mark, Lammle, Todd, and Lammle, Monica
Mastering TCP/IP for NT Server
Sybase books Copyright 1997 SYBEX Inc.
ISBN: 0-7821-213-3

Virtual Private Networking A View From The Trenches
By Bruce Perlmutter with Jonathan Zarkower
Prentice-Hall, Inc Copyright 2000
ISBN: 0-13-020335-1

Security Complete
SYBEX Books Copyright 2001 SYBEX Inc
ISBN: 0-7821-2968-4

MCSE Training Kit Microsoft Windows 2000 Professional
Copyright 2000 by Microsoft Corporation
ISBN: 1-57231-901-1

MCSE Training Kit Microsoft Windows 2000 Server
Copyright 2000 by Microsoft Corporation
ISBN: 1-57231-903-8

Zwicky, Elizabeth D. Cooper, Simon, and Chapman, D. Brent
Building Internet Firewalls 2nd Edition
Copyright 2000 O'Reilly & Associates, Inc
ISBN: 1-56592-871-7

Strebe, Matthew and Perkins, Charles
Firewalls 24 Seven
Copyright 2000 SYBEX Inc
ISBN: 0-7821-2529-8

SANS Security Essentials II: Network Security Overview
Copyright 2003 The SANS Institute

SANS Security Essentials III: Internet Security Technologies
Copyright 2003 The SANS Institute

SANS Security Essentials IV: Secure Communications

SANS Security Essentials V: Windows Security

Appendix A: Sample Policy

Sample Policy- Modified copy of Server Security Policy found on SANS website.
http://www.sans.org/resources/policies/

Server Security Policy

1.0 Purpose
The purpose of this policy is to establish standards for the base configuration of internal server equipment that is owned and/or operated by Goodman Inc. Effective implementation of this policy will minimize unauthorized access to Goodman Inc. proprietary information and technology.

2.0 Scope
This policy applies to server equipment owned and/or operated by Goodman Inc., and to servers registered under any Goodman Inc.-owned internal network domain.

This policy is specifically for equipment on the internal Goodman Inc. network. For secure configuration of equipment external to Goodman Inc. on the DMZ, refer to the *DMZ Security Policy*.

3.0 Policy

3.1 Ownership and Responsibilities
All internal servers deployed at Goodman Inc. must be owned by an operational group that is responsible for system administration. Approved server configuration guides must be established and maintained by each operational group, based on business needs and approved by the security department. Operational groups should monitor configuration compliance and implement an exception policy tailored to their environment. Each operational group must establish a process for changing the configuration guides, which includes review and approval by the security department.

- Servers must be registered within the corporate enterprise management system. At a minimum, the following information is required to positively identify the point of contact:
  - Server contact(s) and location, and a backup contact
  - Hardware and Operating System/Version
  - Main functions and applications, if applicable
- Information in the corporate enterprise management system must be kept up-to-date.
- Configuration changes for production servers must follow the appropriate change management procedures.

3.2 General Configuration Guidelines

- Operating System configuration should be in accordance with approved the security department guidelines.
- Services and applications that will not be used must be disabled where practical.
- Access to services should be logged and/or protected through access-control methods such as TCP Wrappers, if possible.
- The most recent security patches must be installed on the system as soon as practical, the only exception being when immediate application would interfere with business requirements.
- Trust relationships between systems are a security risk, and their use should be avoided. Do not use a trust relationship when some other method of communication will do.
- Always use standard security principles of least required access to perform a function.
- Do not use root when a non-privileged account will do.
- Servers should be physically located in an access-controlled environment.
- All servers should be equipped with redundant or N+1 Power Supplies.
- Servers are specifically prohibited from operating from uncontrolled cubicle areas.

3.3 Monitoring

- All security-related events on critical or sensitive systems must be logged and audit trails saved as follows:
    - All security related logs will be kept online for a minimum of 1 week.
    - Daily incremental tape backups will be retained for at least 1 month.
    - Weekly full tape backups of logs will be retained for at least 1 month.
    - Monthly full backups will be retained for a minimum of 2 years.
- Security-related events will be reported to the security department, who will review logs and report incidents to IT management. Corrective measures will be prescribed as needed. Security-related events include, but are not limited to:
    - Port-scan attacks
    - Evidence of unauthorized access to privileged accounts
    - Anomalous occurrences that are not related to specific applications on the host.

3.4 Compliance

- Audits will be performed on a regular basis by authorized organizations within Goodman Inc.
- Audits will be managed by the internal audit group or the security department, in accordance with the *Audit Policy*. The security department will filter findings not related to a specific operational group and then present the findings to the appropriate support staff for remediation or justification.
- Every effort will be made to prevent audits from causing operational failures or disruptions.

4.0 Enforcement

Any employee found to have violated this policy may be subject to disciplinary action, up to and including termination of employment.

5.0 Definitions

Term	**Definition**
DMZ	De-militariezed Zone. A network segment external to the corporate production network.
Server	For purposes of this policy, a Server is defined as an internal Goodman Inc. Server. Desktop machines and Lab equipment are not relevant to the scope of this policy.

6.0 Revision History