

A Whitepaper on Building and Managing Secure Networks

Prepared for:
Southern Methodist University
School of Engineering and Applied Science, Houston
and
The Association of Information Technology Professionals
Houston Chapter

Written by:
Wes Noonan, MCSE/MCT/CCNA/CCDA/NNCSS
wnoonan@houston.rr.com
WJN Consulting, LLC.
www.wjnconsulting.com

Building and Managing Secure Networks

One of the most difficult tasks for network engineers to perform is building not only secure, but manageable networks. This is largely due to conflicting interests at either end of the spectrum. On one hand we have the user needs, which tend to work more towards the non-secure/unmanaged side of the house. On the other hand we have the need to secure and manage the network. In a perfect world, we would be able to meet both needs all the time. Unfortunately, we aren't in a perfect world. So how does one go about trying to bridge the gap? That's what we are going to discuss.

The first topic to discuss is planning for security. The first thing to remember is that security is a process and a policy, not a product. Don't expect to implement a device, or a piece of software and achieve security. You need to build a structure and you need to have a process and a policy first. Only then can you start implementing the security solutions and products.

The Security Policy

The Security policy is the guidelines that define what how your network is going to function. This policy is a living document that will change as the network grows, but should always remain an up to date reference to eliminate confusion and complications from your network design and implementation. There are 3 areas of implementing a security policy that Cisco has defined for security management. They are *preparation*, *prevention* and *response*.

Preparation

When I was in the Marine Corps, they used to talk about the 7 "P's" – Proper Prior Planning Prevents Piss Poor Performance. This holds true for networking. You need to plan and prepare *before* you implement anything. As they say, "an ounce of prevention is worth a pound of cure". There are a couple of topics that address the preparation phase:

- Create usage policies
- Conduct a risk analysis
- Establish a security team structure

Creating Usage Policies

The first thing to do is set responsibilities. You need to define what the roles and responsibilities are of the users in your network. You need to document what the users can, and cannot, do. Your users need to *understand* what the security policy is, why they should follow it and what their responsibilities for security are. One of the most common means to gain access to a network is to hack the users. By that I mean doing things like calling the users and asking them for their passwords. This is a practice known as social engineering. For example, your users need to understand that no one is going to ask for their password, so if someone does do this, they should not give it out.

You also need to establish acceptable usage guidelines and statements. Quite frankly, this is the time to get the lawyers involved. You need to make sure that you clearly define what can and cannot be done on the network. You need to make it clear whether monitoring is occurring and finally you need to define what punitive actions can and will be taken in the event that security attacks or exceptions to the acceptable use policy are detected. I remember reading a story a few years ago about a hacker who got out of punishment because the system he hacked had a "welcome" banner. The defense amounted to "he was invited into the system so it's not his fault". It actually won!! Make sure your usage guidelines are clear, published and when relevant, that they require acknowledgement or a signature to allow for access to be permitted. Leave the "politeness" for other times. Don't welcome people into your Internet gateway. ☺

You also need to create an administrative acceptable use statement that details the procedures for user account management, policy enforcement and privilege/rights assignment. I remember a common problem in Windows NT networks in regards to group memberships. If you create a global group while highlighting a user, the admin tool assumes that user should be a member of the group in question. Many administrators were unaware of this, and since the tool defaulted to highlighting the first user in the list when opened, I used to spend quite a bit of time at customers cleaning out the user whose name started with "A" out of groups. Procedures that explained the process would have gone far in those environments. You need to

Written by Wes Noonan, MCSE/MCT/CCNA/CCDA/NNCSS
wnoonan@houston.rr.com
WJN Consulting, LLC.

Building and Managing Secure Networks

document things like password policies and access times. Make sure though that all of your policies present a single uniform statement. Nothing is more confusing than conflicting instructions.

Finally, ensure that the usage policies can be followed. A usage policy that is a hindrance to work is not an effective policy. If the policy starts to get in the way of doing the job, and thus stops being followed, it isn't doing any good anymore and needs to be re-evaluated.

Conduct a Risk Analysis

You need to identify where your risks are, risks to the network, to the resources and to the data. Risk analysis doesn't mean you need to be a hacker though. You don't need to identify every single possible method of attack that there is. You do need to identify the portions of the network that exist, assign a threat rating to each portion and then apply a level of security to address the threat level. Many security administrators fail to grasp this concept. Not everything needs to be defended like Fort Knox.

The most basic method of risk analysis is to assign one of three levels of risk to your network resources:

- Low Risk – These are systems or data that if compromised would not disrupt the business, cause legal issues or result in any financial ramifications. The targeted system can be easily rebuilt and cannot be leveraged to provide further access to systems. In many cases, contrary to popular opinion, devices like bastion hosts could well fall into this category.
- Medium Risk – These are systems that if compromised would cause a moderate disruption to the business, minor legal or financial ramifications or could be used to provide access to other resources. These systems would require moderate effort to restore, and the restoration too could be disruptive to the system.
- High Risk – These are systems that if compromised could cause significant disruption to the business and could result in significant legal or financial ramifications. In extreme cases, they could even threaten the safety of personnel. These systems would take significant effort to restore, at a significant disruption and cost to the business.

Based on the three risk levels, you then want to assign the appropriate level of risk to the various network systems – core network devices, distribution network devices, access network devices, network monitoring devices, network security devices, email systems, network file and print servers, network applications servers (i.e. DNS or DHCP), data application servers (i.e. Oracle or SQL), desktop computers, and other network devices not already covered.

When you go to assign risk levels, make sure you think outside of the box. For example, what is the risk level of a DHCP server being compromised? Not high you say? Consider that it contains a mapping of every single in use IP address and MAC address on your network. Sounds like spoofing made easy in the wrong hands to me.

Once you have assigned the risk levels, the next thing to do is identify the types of users you have, and then to assign what privileges they will have on each system. There are generally 5 types of users:

- Administrators – These are **trusted** internal users responsible for network access. I remember a customer one time asking what could be done to restrict what the administrators could do on the network. I replied better hiring practices. These are your administrators. If you don't trust them, they shouldn't be your administrators.
- Privileged Internal Users – these are users with a need for greater access, for example technicians and help desk staff.
- Users – these are your general access users.
- Partners – External users with a need to access some resources. This may be access to ERP systems, web sites or other data. This could also be vendors who provide support.
- Others – This is everyone else.

Building and Managing Secure Networks

While every network is different, and every administrator needs to define the risk levels of their individual network, Cisco provides a quick reference starting guide to define how to assign some of the classifications we have discussed. Remember that the following table isn't an absolute statement, but is a good starting point.

System	Description	Risk Level	Types of Users
ATM switches	Core network device	High	Administrators for device configuration (support staff only); All others for use as a transport
Network routers	Distribution network device	High	Administrators for device configuration (support staff only); All others for use as a transport
Closet switches	Access network device	Medium	Administrators for device configuration (support staff only); All others for use as a transport
ISDN or dial up servers	Access network device	Medium	Administrators for device configuration (support staff only); Partners and privileged users for special access
Firewall	Access network device	High	Administrators for device configuration (support staff only); All others for use as a transport
DNS and DHCP servers	Network applications	Medium	Administrators for configuration; General and privileged users for use
External e-mail server	Network application	Low	Administrators for configuration; All others for mail transport between the Internet and the internal mail server
Internal e-mail server	Network application	Medium	Administrators for configuration; All other internal users for use
Oracle database	Network application	Medium or High	Administrators for system administration; Privileged users for data updates; General users for data access; All others for partial data access

Establish a Security Team Structure

The last part of preparation is to establish a functional security led by a team manager. You should include personnel from other areas of your company. These folks need to be well versed in not only the policies, but they need to be practiced in their skills and they need to be prepared to execute the response to a situation that may occur. This is often going to require additional training of these people. There are three main areas of responsibility for your security team:

Policy Development – These folks need to be establishing and reviewing the security policies discussed earlier. This should occur on an annual basis at a minimum.

Practice – This is the time where your security team is conducting risk analysis, addressing change control requests, reviewing security alerts from vendors as well as mailing lists (i.e. CERT – www.cert.org or Security Focus – www.securityfocus.net) and turning security policy requirements into technical implementations. This is where the planning starts getting put into practice (bad pun intended) ☺.

Response – I rode motorcycles for a while and we had a saying. “There are two types of riders, those who have hit the ground and those who will”. Security is like that as well. You can have the best plans and practice and you still might get compromised. In that event, your team members need to know what their responsibility is in defining, troubleshooting and fixing a security compromise. Each team member must know all of the security features of the equipment and applications that they are responsible for.

Building and Managing Secure Networks

Prevention

Prevention falls into two categories: approving security changes and monitoring security of your network.

Approving Security Changes

Security changes are nothing more than changes to network equipment that could have an impact on the overall security of the network. Your security policy needs to define, in plain English, what your security configuration requirements are. In short, keep the technobabble to a minimum and instead be concise and clear in your statements. Every administrator needs to define a unique set of requirements for their individual network and organization. Simply put, you need to use and follow a change control policy. I went to a customer site one time and they asked us if we had software that could help them manage their firewalls. It seemed that firewalls went down all the time. Upon further investigation, we discovered that the problem was that there were 3 different NOCs and each NOC would change the firewall policies without telling anyone. In the end, we broke our salesman's heart when we told the customer that all the software in the world wouldn't help this problem. This was a people policy issue that would only be addressed by implementing a strict change control policy.

To create your change control policy, your security team needs to define the requirements needed to identify specific network configuration and design issues that meet those requirements. In short, define what you need so that you know when you need it. It is recommended that the security team review the following changes *before* anything is changed:

- Any firewall configuration changes.
- Any Access Control List (ACL) changes.
- Any Simple Network Management Protocol (SNMP) configuration changes.
- Any changes or updates in software that differs from the approved software revision list.

Additionally, the following guidelines should be followed in terms of change control practices:

- Change passwords on a routine basis, especially when someone who knew the password was leaving. I worked at a bank where when a person left, while they were saying goodbye to admin A, admin B was disabling their accounts. Harsh, but a good policy.
- Restrict access to network devices to approved users only.
- Audit and ensure that the correct software revisions are running on your equipment.

Monitoring Security of Your Network

Monitoring the security of your network is similar in concept to network monitoring, however the focus of security monitoring is to look for changes that indicate that a security violation is occurring. In order to determine if a security violation is taking place, you must first define what a security violation is. In risk analysis we defined what the level of monitoring we needed was based on the threat to the system. In approving security changes, we identified specific threats to the network. Understanding this allows us to define what we need to monitor and how often we need to monitor to identify violations.

As an example, a firewall is considered a high-risk network device. As a result, firewalls need to be monitored in real time. We also know that we should be monitoring for any changes to the firewall configuration. To that end, a good security practice would then be to monitor the firewall to report things such as failed login attempts, unusual traffic, changes to the firewall, access granted to the firewall and connection setups through the firewall.

Following this example, it is necessary to create monitoring policies for each area identified in your risk analysis. As a rule of thumb, you could monitor low risk devices weekly, medium risk devices daily and high-risk devices hourly. This is just a rule of thumb however and if you think you need more rapid detection simply decrease the time frame to meet your needs.

There are a number of tools that you can use to monitor your firewalls and networks for security violations. BMC Software (www.bmc.com) makes PATROL Agents that can monitor Checkpoint and PIX firewalls,

Written by Wes Noonan, MCSE/MCT/CCNA/CCDA/NNCSS
wnoonan@houston.rr.com
WJN Consulting, LLC.

Building and Managing Secure Networks

Security Management with INCONTROL and Network and Traffic Management with PATROL DashBoard and PATROL Visualis, HP (www.hp.com) makes HP OpenView, NetIQ (www.netiq.com) makes End2End Performance Monitor Suite, Chariot and the WebTrends Firewall Suite, NetScout (www.netscout.com), makes nGenius, MicroMuse (www.micromuse.com) makes NetCool and Tivoli (www.tivoli.com) makes NetView and SecureWay.

Another set of tools for monitoring the security of your network that are worth looking at are Intrusion Detection Systems (IDS). These are tools that can look at the traffic on a network or connecting to a host and analyze the traffic for signs of intrusion and misuse. IDS systems generally fall into 2 categories, host based and network based. Host based IDS systems are installed on the hosts that you want to monitor and screen the traffic connecting to the host to ensure it's validity. Some examples of host based IDS systems are NetworkICE/ISS (www.networkice.com) BlackICE Defender and RealSecure, and Axcent/Symantec (www.symantec.com) Intruder Alert. Network based IDS systems function by being located in the network in locations that allow them to view the traffic passing over a segment. Some examples of network based IDS systems are Cisco (www.cisco.com) with Cisco Secure IDS, ISS (www.iss.net) with RealSecure, Network Flight Recorder (www.nfr.net) with NFR NID and TripWire (www.tripwiresecurity.com) with TripWire. Common locations for network based IDS systems to be located are:

- On internal segments
- In front of a firewall
- Behind a firewall
- Behind dial-up/RAS server
- At extranet connections

Another real good website for IDS information is www.icsalabs.com. They perform vendor independent certification of IDS systems and have a really good whitepaper that introduces IDS systems in more detail than we have time for. One thing to keep in mind with IDS systems though is regardless of how good they sound, they are not silver bullets and they cannot replace the need for human beings to do something productive with the information they collect.

The last thing that your security monitoring policy needs to define is how to notify the security team that a violation has occurred. The odds are that a well implemented system of firewall and network monitoring as well as IDS should detect violations well before you the administrator ever do. To this end you need to make sure that your policy defines a system of triggering notifications through an operation center usually using SNMP, and then to the security team using pagers or email. An important thing to make sure of is that you reduce the amount of false positive pages by making sure your software isn't configured to trigger too often, but also doesn't trigger too late. It's a fine line to walk, but a security team that gets numerous pages when nothing is really wrong quickly becomes inefficient due to the "boy who cried wolf" syndrome. In these cases, when a valid page comes in, the administrator discounts it as "probably not a real problem". Not a good place to be.

Response

Response can be broken into 3 parts, *security violations, restoration and review*.

Security Violations

Practice makes perfect and planning how to respond to a security violation will make your response much more effective. Let's face it, when it has hit the fan things are hectic enough that the last thing you need is nothing to guide you in your actions. By making the decisions on how to react to security violations *before* a violation has occurred, you are going to make it that much easier to deal with when it happens. In a sense, having the plan ahead of time is like doing fire drills. It makes it that much easier to find the door in all the smoke.

The first thing that needs to happen when a violation is detected is to notify the security team. You need to define a 24/7 policy on exactly how to react to an issue. I know of a number of companies that got hit

Building and Managing Secure Networks

extremely hard due to Code Red II because it hit Friday night and by the time their IT staff found out Saturday, Sunday or even Monday morning the damage had already been done.

The next topic to address is the level of authority that the security team has to make changes, and in what order the changes should be made. Some examples of corrective actions that can be taken are:

- Implementing changes to prevent further access to the violation.
- Isolating/disconnecting the compromised systems or the source.
- Contacting the provider to begin tracing the attack.
- Using recording devices, logs and sniffers to gather evidence.
- Contacting appropriate authorities.
- Restoring systems according to a prioritized list.
- Notifying internal management and legal personnel.

Additionally, you need to ensure that any changes made to the systems that do not have management approval are documented for future reference.

Some people ask why they need to collect and maintain the information during a security violation. They are so busy trying to deal with the problem now, they lose sight of understanding what happened, and why it happened, so they can keep it from occurring again. There are two very big reasons for collecting this data for future reference: to determine the extent to which systems have been compromised by a security attack and to prosecute violations.

To determine the extent the extent of the violation, you need to perform the following:

- Record the event by obtaining sniffer traces of the network, copies of log files, active use accounts and network connections. Tools such as BMC Software PATROL Visualis which can record and playback network traffic are particularly valuable in these kinds of “forensic” tasks since they can actually show you the traffic patterns as they occurred, providing information such as the point of origin.
- Disable any accounts used, disconnect the equipment from the network and if need be disconnect from the Internet to limit any further compromises. Once Code Red II took hold, this was the only effective measure in preventing further spread.
- Backup the compromised system to aid in a detailed analysis of the damage and attack methods that were used.
- Look for other signs of compromise. One of the best responses to a security breach I have ever seen was a gentleman who responded to the question “why should I rebuild the server? I know this exploit doesn’t place any Trojans on my system” with “If your system was able to be compromised by this, what makes you think it *wasn’t* compromised by a Trojan at some other time”. It is very common for a system to be compromised at multiple levels, if it is compromised at all. Heck, Nimda exploited no less than 15 different holes and breaches.
- Maintain and review security and device logs as they often will provide clues as to the method and nature of the attack. This is easily one of the least done tasks that I witness take place. People setup syslog servers, they enable auditing and logging on their servers but they never go back and check the logs until after a problem has occurred. This is nice, but it is too late. In many cases, since people never look at the logs on a routine basis, they have no idea what is a normal event versus what is an abnormal event. This makes the logs effectively useless in many cases. In fact, I worked at a company whose IIS logs were so large that they couldn’t even be loaded so that they could be read. This is not exactly a case of effective logging. In fact, in that situation you might as well not log at all.

If you want to prosecute, or take legal action, then you need to have your legal department work with you to review the process of gathering evidence and the involvement of the authorities. You may also need to involve human resources in the event that the violation was internal in nature. The key thing to remember

Building and Managing Secure Networks

here is that you are not a lawyer, nor are you an HR rep. Tread lightly, and let them do their jobs in these cases.

Restoration

The ultimate goal of your security policy and processes should be the restoration of services. You need to define and provide information regarding the backup and restoration procedures that will be used.

Review

The review process is the final step in creating and maintaining a security policy. There are three things to review: policy, posture and practice.

The security policy is a living document that needs to change when required. As best practices change, your policy needs to be checked against the best practices to make sure it is up to date. You should review your policy against both vendor and CERT (www.cert.org) resources for tips, practices, and security updates and alerts that you can incorporate into your policy.

You also need to compare the posture of your network against the desired security posture. This is where outside auditing comes in. Because you know your network, quite frankly you are the last person that should be validating its posture. It is generally recommended to have an outside company audit your network and security practices at least once a year. Many companies schedule this at the same time they perform their fiscal audits, and as a result many of the “big 5” consulting companies have IT organizations as counterparts to their line of business operations.

Finally, practice the policies you put into place. You should schedule and perform drills to test the security staff and ensure that everyone knows what their responsibilities are and what they need to do when a violation is detected. While having notice of a drill is worthwhile, it is also important to have unannounced drills periodically. It’s kind of like the part in “Heartbreak Ridge” where the troops are complaining that they always ambush the other group right here when Clint Eastwood says “kind of makes it easy to get out of an ambush when you know when and where it is happening”. Apply the same philosophy to your security policy testing. It is also common to perform the drill in conjunction with your posture testing to further assist in illustrating gaps in the process and procedures that need to be addressed and corrected.

With the security policy defined and in place, let’s discuss some common network designs that can provide for more security.

The Secure Design

The key to designing a secure network is to implement the security policy you developed, not to purchase a product. One of the great fallacies of network security is that some device – a firewall, IDS, port filter, NAT router, etc. will provide the level of security that you want. The reality is far from that. Like I mentioned, security is a philosophy or a process, not a “thing”. The first thing to understand in designing for security is who the enemies are.

The Enemies

There are 4 main categories of people who threaten security on a network, hackers, unaware staff, disgruntled staff and snoops.

Hackers – This has become a generic catch all term to define anyone who tries to gain access to other peoples systems. While the vast majority are simply content to prove that they could get in, there are a rare few known as “crackers” who are more malicious and desire to cause as much damage and headaches as they can.

Unaware Staff – These are the unwitting victims that folks like the hackers love. They choose poor passwords, run every executable that is sent to them, and disregard many security policies because “that’s not my job”. Unfortunately, these folk contribute as much to security breaches, if not more, than any other

Building and Managing Secure Networks

group. The best that you can do is implement security policies that force these users to adhere to the policies set forth. After all, if you don't save them from themselves, who will? ☺

Disgruntled Employees – These folks have the potential to cause more damage than any other group. This is due largely to the fact that they are considered “trusted” by most organizations, and in some cases that view isn't changed until it is too late. Everyone knows the stories. The admin who changed all the passwords and didn't tell anyone, or the user that deleted all their files before they left. Sometimes the best safeguards against these types of users is simply having good recovery and backup procedures.

Snoops – This is your classic category of curious user. The employee who is taking a computer class and wants to put some knowledge to work, or the outside user who happened to find an open door that they just can't help but walk through. In some cases the risk may be as high as espionage, in other cases the risk is simply a nosy person looking at or doing things they shouldn't

The Risk

There are 7 big risks to address with these enemies.

Viruses – Viruses have become the single most prevalent threat to security today. Some viruses are relatively harmless, but others destroy data and restrict access.

Trojan Horse Programs – These programs may appear as useful software programs and games, however they have an ulterior motive. They can install on a system and perform tasks ranging from spreading themselves to other systems, deleting data or in many cases they can provide remote access to the system for the originator of the Trojan, for example Back Orifice.

Vandalism – This risk is targeted primarily at web sites where the person vandalizing, sometimes called tagging, seeks to make a statement. Sometimes it is as simple as “you've been tagged” and other times it is used to make political statements.

Attacks – There are innumerable types of network attacks, however they generally fall into three classifications: reconnaissance attacks, access attacks and denial of service (DoS) attacks.

- **Reconnaissance Attacks** – These attacks are generally used as information gathering where by hackers collect data which can be used in the future to later compromise networks. Examples of tools used for this are sniffers and scanners.
- **Access Attacks** – These attacks are generally conducted to attempt to gain access to systems, particularly file and authentication services. Examples of tools used here are things like password crackers and L0phtCrack.
- **DoS Attacks** – These attacks are generally used to prevent access to part or all of a computer system. The general function is to send unmanageable data to a machine in such volume that it is unable to respond to legitimate requests in a timely fashion. A more malicious form of DoS is the Distributed Denial of Service (DDoS) attack, which has gotten much press in recent months.

Data Interception – Any data sent on a network can be intercepted by unauthorized people. The purposes might be to simply listen in, to hope to capture passwords off the wire, or even to change the data that is being transported. This is the arena that things like IP Spoofing – acting like a different IP address, and TCP Hijacking – Attempting to insert ones self into the conversation by guessing the next TCP sequence being used come into play. This is also where significant data loss can occur. Credit card fraud and things of that nature come into play here.

Social Engineering – This is the increasingly prevalent act of gaining access to systems through non-technical means. The classic example is to call a user and act like you are part of technical support and tell them that you need their password. You would be amazed at how often this works. I worked at a company that I tried this on the first day of work, before anyone knew who I was. I got 100% of the passwords that I asked for.

Building and Managing Secure Networks

Spam – No, not the lovable luncheon meat of Monty Python fame, but the term used to refer to unsolicited commercial email (UCE). Spam risks fall into two distinct categories. The first is receiving it, which consumes time, bandwidth and resources to deal with and delete. The second is sending it, for example by leaving your mail relay open. Spammers look for open relays and then use your relay to send their spam, taking up your bandwidth instead of theirs for the job.

The Fix

Now that we know the enemies, and now that we know the risk, what can we do about it? There are number of tools at our disposal that can assist us in this effort ranging from Anti-Virus packages, hardware, firewalls, IDS and education. The most important thing to remember though is that all the technology in the world isn't a substitute for good security policies that are followed by personnel.

Anti-virus Software – This is about the only thing that can be done to effectively mitigate damages from viruses. A good anti-virus package should be able to detect and stop all existing viruses, as well as have frequent and timely updates available for download from their website. You should keep the software updated on a regular basis, and should perform weekly scans on your systems at a minimum. Personally, I recommend nightly updates and daily scans. To support such an aggressive policy, it is important that your anti-virus software support user customization on what and where to get the updates. I recommend Network Associates Virus Scan or Symantec Virus Scan (Corporate Edition only).

Access Control – Access control servers are like the gatekeepers to your network. By using access control servers, you force users to present some form of validation before they are granted access to your systems. Access control spans topics from network logon and authentication to RADIUS Servers to TACACS+ servers to Access Control Lists (ACLs) all of which control what a user can and can not access to one degree or another. Access control servers are like the gatekeeper handing out badges to the people who need them. Badges!!! Badges!!! We don't need no stinkin' badges!!! (sorry, couldn't help it) ☺

Firewalls – Firewalls have earned a reputation as a panacea of sorts. Like most panacea's, the reality is that they can be much closer to snake oil than to a cure-all for your network woes. All of this is not to say that firewalls are useless however, they just aren't the *only* things. Firewalls function by restricting access to network resources as well as enforcing security policies. A firewall is similar to a door lock in the sense that without the key, entry is denied. Firewalls don't just stop at securing the entry to the network though; they can also run filters and proxies that can actually filter specific data based on content as well.

Encryption – Encryption ensures that the data being transmitted cannot be deciphered by anyone other than the authorized destination. Encryption has become more prevalent as more people have begun using the Internet to transmit data. Encryption is what has allowed people to use Virtual Private Networks (VPN) to connect sites in a secure fashion. VPN's function by creating a "tunnel" through which the data is transported, making the interception and reading of the data virtually impossible to anyone other than the intended destination. VPN's are analogous to armored cars in their function, protecting sensitive data in an open and hostile world.

Intrusion Detection – If firewalls are like the locks on the doors and VPN's are like the armored cars transporting valuables in a hostile world, intrusion detection is like the surveillance camera recording everything that passes by. A good IDS provides a round the clock surveillance of the network, analyzing packets and looking for unauthorized activity. When this activity is detected, the IDS can send alerts to management consoles that in turn can isolate and stop the unauthorized traffic. Most good IDS's were able to detect Code Red and prevent it from ever entering a network.

Network Scanning – Network scanning is the method by which you the administrator can check your systems and ensure that any vulnerability that might exist get caught and stopped, as well as of course any hacker scanning to find these same vulnerabilities. Network scanning can be as simple as basic port scanning using tools like Nmap (www.insecure.org) to as advanced as using vulnerability scanners like Nessus (www.nessus.org). The goal of scanning is to make sure everything you think is protected really is. In a sense, it's like walking around and making sure the doors are all locked and the windows are all closed.

Written by Wes Noonan, MCSE/MCT/CCNA/CCDA/NNCSS
wnoonan@houston.rr.com
WJN Consulting, LLC.

Building and Managing Secure Networks

Updates – One of the most overlooked and important fixes is the timely application of updates. The best case in point is the Code Red virus. Microsoft had a hotfix for Code Red almost 2 months *before* Code Red was released. Code Red was 100% preventable, yet so many companies and users had not applied a fix that was 2 months old that Code Red was still able to devastate numerous networks and the Internet in general. Subscribe to vendor mailing lists, or to the venerable “Bug Traq” mailing lists at www.securityfocus.net to ensure that you have timely notification of critical updates. If you run Microsoft servers, make sure your desktop runs the Critical Update Notification tool so that you can get notification of patches as Microsoft releases them with little to no effort on your part to track them down.

Expertise – The last, and the most important fix, is technical expertise. The best tools in the world are worthless in the hands of someone who doesn't know how to use them. If you were expecting software and hardware to be the solution to your security problems, you would be putting your faith in the wrong things. You need good people to compliment the tools to really have a winning combination. Another level of expertise that you should leverage is outside expertise. The reality is that as network admins, we aren't hackers and most of us don't have time to be hackers. However, there are companies who do have hackers, and while it may sound unsavory, it is a good practice to bring in the people with that expertise to perform a walkthrough and sanity check of your environment to make sure everything is good to go. This doesn't mean that your staff security folks can't do the job, it simply means that in many cases you have other things your staff guys need to worry about. In addition, a fresh pair of eyes is more likely to discover problems that people familiar with the situation might have inadvertently overlooked.

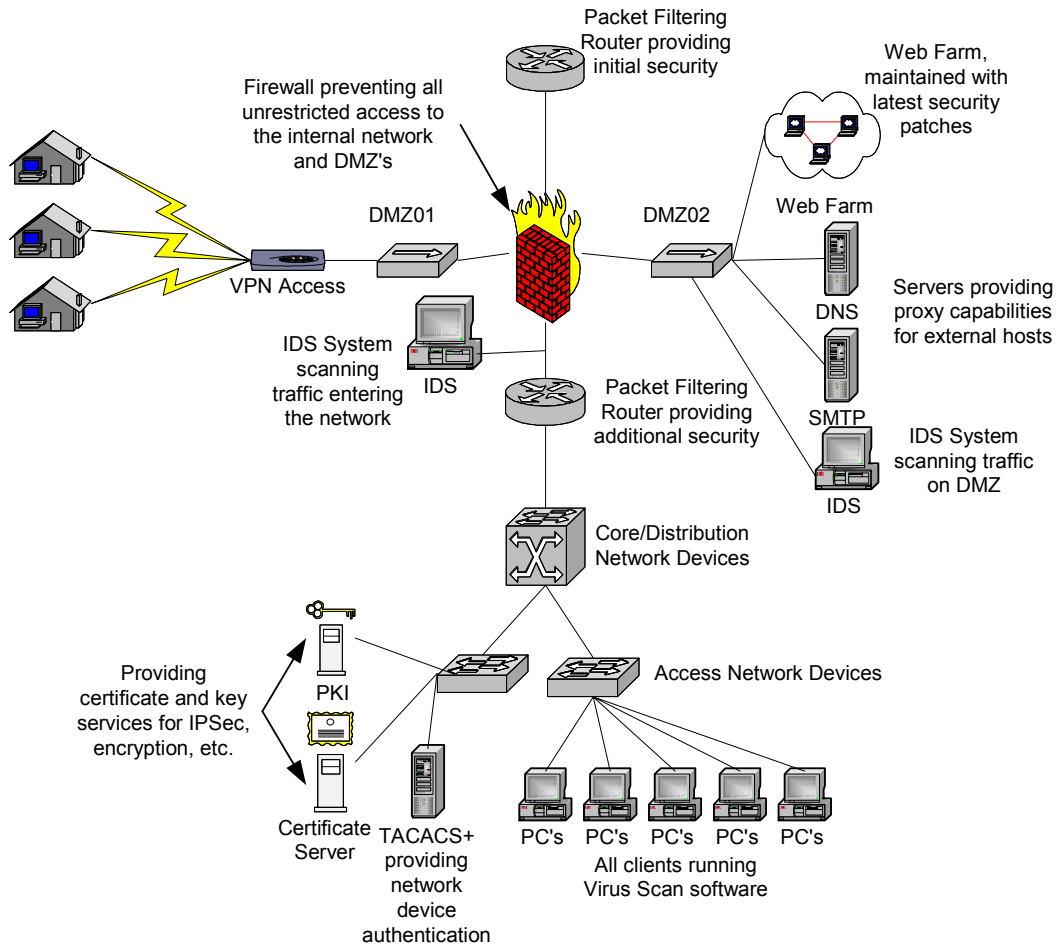
Summary

So how can we build secure and manageable networks? Two words - planning and implementation. Sounds pretty simple, doesn't it? There are two main things to do to ensure that your network is as secure and manageable as possible. The first thing is to establish a good security policy. Prepare for security by creating usage policies, conducting risk analysis and establishing a security team. Next, you need to look at prevention by setting up a change control process and by establishing a monitoring policy. The last component of a good security policy is to define the response that will be taken. This is broken down into responding to security violations as well as defining the methods of restoration and after action review of the situation. In reading this, and sitting through the seminar, you might be a little disappointed that more time wasn't devoted to firewall positioning, IDS, etc. The reality is that most company's security mistakes are made before the equipment is purchased. Planning and policies are the rather inglorious part of the job, and as a result are often overlooked due to the desire to “do something”. Without good planning, the best implementation is likely flawed before the sun sets on its first day. Only after the proper planning has gone into place should the design and implementation be looked at and acted upon. A good design requires an understanding of the enemies, the risk and the fixes. As much as we all worry about the bored 14 year old hacker, realize that internal employees are the biggest enemies out there. Once you have a grip on the enemies, the next task at hand is to consider the risks and their fixes. One of the biggest, and easiest, risks to deal with today is a virus. Use a good Virus scan product with a good update/scan policy. A product that supports centralized downloads, updates and maintenance is the best product to choose. Consider the risk of attacks, access attacks and DDoS attacks, and implement firewalls and strategic partnerships with upstream vendors to address them. Utilize an IDS to determine the nature of the traffic accessing your vulnerable systems. Like I mentioned, most major IDS systems were able to detect Code Red which, had they been implemented well, would have greatly mitigated the risk that Code Red imposed. Keep on top of vendor updates and apply those updates in a timely and efficient manner. We all joke about the amount of “hotfixes” released by some companies, but always remember that there is a reason why those hotfixes keep getting released. Finally, keep expertise close at hand. Bring in outside auditors, and keep your internal staff up to date. Read the various trade periodicals, RFC's and whitepapers that get released. I frequently get asked how I got some knowledge I have or where I learned something, as if it is a mystery. The sad truth is it comes from being a voracious reader of every technical article and whitepaper that I can get my hands on. I think that a certain degree of geek factor is really required to stay on top of this game, and this should be encouraged to maintain the level of expertise to stay ahead of the curve.

Building and Managing Secure Networks

The most important thing to take away from this paper though is to plan and adhere to a policy. The details of what to implement will come easy, once the details of how to do it are worked out. Security isn't a device, it's a policy and it's a philosophy. Build a good policy and by extension, you will be much more likely to have good security.

A Secure Network Diagram Example



Building and Managing Secure Networks

References

The following articles and Whitepapers were used as references for this presentation:

Cisco Systems (www.cisco.com)

“A Beginners Guide to Network Security”

“Internetworking Technology Overview, Chapter 47, Security Technologies”

“Increasing Security on IP Networks”

“Internetworking Technology Overview, Chapter 6, Network Management Basics”

“Configuration Management: Best Practices Whitepaper”

“Disaster Recovery: Best Practices Whitepaper”

“Capacity and Performance Management: Best Practices White Paper”

“Service Level Management: Best Practices Whitepaper”

“Network Security Policy: Best Practices Whitepaper”

“Network Management System: Best Practices Whitepaper”

“Performance Management: Best Practices Whitepaper”

“Change Management: Best Practices Whitepaper”

Internet Security Systems (www.iss.net)

“A Vision for Complete Protection”

“Creating, Implementing and Managing the Information Security Lifecycle”

“Security Convergence Solutions”

TruSecure Corporation (www.trusecure.net)

“ICSA 3rd Annual Firewall Buyers Guide”

“An Introduction to Intrusion Detection and Assessment”

National Security Agency (<http://nsa1.www.conxion.com/>)

Security Recommendation Guides – Windows 2000 Guides

Security Recommendation Guides – Cisco Router Guides

Tools and Software

Firewalls:

Cisco PIX and IOS Firewall (<http://www.cisco.com/warp/public/44/jump/secure.shtml>)

Firewall-1 (www.checkpoint.com)

Symantec Enterprise Firewall (www.symantec.com)

ISA Server (www.microsoft.com)

IDS:

Cisco Secure Intrusion Detection System

(<http://www.cisco.com/warp/public/44/jump/secure.shtml>)

RealSecure and BlackICE (www.iss.net)

Netprowler (www.symantec.com)

Port Scanners, Penetration Testers and Vulnerability Assessment Tools:

Nmap (www.insecure.org/nmap/index.html)

Nessus (www.nessus.org)

Retina (www.eeye.com)

Internet Scanner (www.iss.net)

Network Monitoring, Performance Monitoring and Fault Management Software:

PATROL 2000, PATROL DashBoard, PATROL Visualis and INCONTROL (www.bmc.com)

HP OpenView (www.hp.com)

WhatsUp Gold (www.ipswitch.com)

Unicenter TNG (www.ca.com)

nGenius (www.netscout.com)

NetCool (www.micromuse.com)

WebTrends Firewall Suite, Chariot and End2End Performance Monitor Suite (www.netiq.com)

Written by Wes Noonan, MCSE/MCT/CCNA/CCDA/NNCSS

wnoonan@houston.rr.com

WJN Consulting, LLC.