Global Information Assurance Certification Paper

But I have a firewall, my network's secure!

An Introduction to network security.

Derran Guinan

October 4, 2004

GIAC Level One Security Essentials Practical Assignment
Version 1.4b – Option 1

Table of Contents

1. Abstract

Network security has never been more complex then today and the trend is that things are getting more complicated, not less. Far too often network administrators and users put all their security faith into perimeter security (firewalls), but perimeter security is only one component of a network security solution. Network security also needs to include intrusion detection, patch management, virus management, physical security of your equipment, and most importantly it must be driven by security policies. There's no question that having a firewall or a patch management solution in place improves network security, but without addressing the other above-mentioned areas your network is not as secure as you thought or hoped. As each category could be a research topic in itself, this paper is only intended to introduce network security and its various components to new network administrators and security officers.

2. Security Policies

2.1 What is a Security Policy?

All organizations today share one common problem, security threats. These threats range from computer hacking and viruses to denial of service attacks.[1] With a virus everyone knows that you should have an anti-virus program to protect your computer or network, but the first step should not be what do I need to stop viruses, but what needs to be protected. This is where a security policy comes in. The policy is there to identify what needs to be protected, the risk of that need being attacked, and the prioritization of that need for protection. [1] As mentioned above there are many facets to network security, but they should all be based on a security policy. Without such a policy, certain security precautions might be occurring but they are most likely occurring on an ad-hoc fashion, and inconsistently across the organization.

2.2 Having the Right Players Involved

The first step in developing a security policy is having a person own and implement the proposed policy. If the organization had a Chief Security Officer (CSO) this would be their job, but in most organizations they will be lucky if they even have a Security Officer. Therefore, the main prerequisite is that the person must be committed to the task, and they must have enough authority to ensure that the policy is completed on schedule.

The second step is to have a Security Committee formed with membership from across the organization, encompassing different business groups.[1] The goal of this committee will be to ensure that the security policy or policies meet the needs of the organization. Since these committee members are from various business groups, at times, they will help the Security Officer understand various business processes and at other times provide a safety check that the proposed policy can be realistically adopted within the organization. There's no point

working on, or passing a security policy that will not realistically work in the organization.

2.3 Understanding the Business Goals

The third step is to understand the business, and how that business generates revenue (if it's a private business). This would also include the direction and goals of the organization as well as ensuring that any policy that is created meets with the organizations current policies, rules, regulations, and any laws that the organization must comply with. To accomplish this, the Security Officer would interview senior management, and business managers to identify business processes within the organization. After these processes have been identified, another list of all of the business assets that are used by the business processes is created. This second list would include physical devices such as routers, switches, hubs, and servers. The size of the list is unimportant (it will be large), but ensuring that all assets that are used by a business process is critical. Missing a process or asset would mean leaving a potential security hole. To assist in determining the assets importance, it's important to attach a value to each asset. This value will allow you to rank your assets level of importance. For example, it's typically more important to ensure that a $500,000 server farm is protected, than a single $5000 server.

2.4 Risk Assessment

Once you have a classification of your assets completed, it's time to begin a risk assessment. To start a proper risk assessment vulnerabilities for each business asset must be identified. This can be done by using a tool to scan your systems and produce a report identifying any known problems. These scanning tools are freely available from the Internet. Even though these tools are very good at detecting vulnerabilities, a manual review must be done to identify any missed problems, or for any systems that a vulnerability tester cannot test.[2] It's also important not to forget about physical resources, such as access into buildings, and server rooms.

Through the risk assessment it is important to raise the awareness among senior management of the potential business costs of security breaches.

2.5 Implementing the Policy

Once a proper risk assessment is complete the Security Officer and the Security Committee can review the assessment and determine how to proceed. At this time, the committee may choose to handle various risks as separate policies, or they may choose to develop all the potential risks as one security policy. How the risks are presented is very important, and can be very strategic both in regards to technology and organizational politics. For example, breaking the policy up into smaller policies is advisable when you have some 'low hanging

fruit', things that can be easily fixed. These quick fixes are always good, because they show progress which management likes to see. Additionally, sometimes presenting one large security policy for all of the potential risks can take months to approve due to the complexity in understanding all the implications around it.

2.6 Auditing and Monitoring of the Policy

Once the policy has been approved and implemented, ensuring that the policy is being adhered to can be just as time consuming and difficult as the development. For that reason periodical audits by internal auditors (if the organization is large enough to have their own) or an outside auditing firm is necessary, and any deviation from the policy should be reviewed, a reason for the deviation determined and assuming there's no reason for the deviation the individuals should be disciplined. [2] If policies are not enforced strictly it sends the wrong message to the organization, and the effectiveness of the policy(s) is lost.

Audits and constant monitoring of policies go hand in hand. As technology changes so will your policies. A good security policy is intended to be a work in progress, not some policy that gets thrown on the shelf and never gets updated, and therefore eventually loses its effectiveness.

3. Perimeter Security

3.1 Defining Perimeter Security and Firewalls

What is perimeter security? Is it a firewall, or is it more? The answer is yes on both counts. Perimeter security is the protection of your logical network (Subnet, or multiple subnets) from the Internet. How do you protect your logical network? Firewalls. A firewall is simply a device that decides what traffic should be allowed in or out of your network based on a predefined set of rules. The rules for a firewall typically come from the organization's security policy, which as described above are determined through identifying potential security risks.

3.2 Where should firewalls be located?

To start with, firewalls should be located between your network and any other external network, such as the Internet. This spot is what D. Brent Chapman and Elizabeth D. Zwicky call a 'chokepoint' in their book, Building Internet Firewalls (O'Reilly, 1995). It's termed a 'chokepoint, because it's the single point between two or more networks, through which all traffic must pass. By placing a firewall at this location(s), you are now able to control, monitor and log traffic to and from your network.[3]

3.3 Is the Perimeter Changing?

As technology is continually changing so is the perimeter. Every network still needs a firewall to protect the network, but the question is now how many and where should they be located.[4] A firewall is still needed at the 'chokepoint', but additional firewalls or zoning of the one firewall is needed to separate the various types of servers such as database, infrastructure and any web servers from the workstations, the public and each other. By locating the various servers in different zones, and keeping workstations in another, completely separate zone you can now apply different security policies against each zone.[5] For example, all the database servers should not be accessible to and from the public, while all the web servers, or Citrix servers will be (they would be located in what's called a demilitarized zone or DMZ). The main premise to follow is that of 'least privilege'. Users, servers, workstations, administrators, anyone and anything should only have the privileges that they need to perform their function.

It's also important to remember that malicious activity does not only occur from outside of the organization. Organizations must still protect themselves from internal threats from disgruntled employees to network snoopers.

3.4 Types of firewalls

Typically, firewalls are categorized into one of these categories:

- Packet Filters
- Circuit-level Gateways
- Application Gateways
- State-full inspection

Packet Filters

First off, what is a packet? Basically, a packet is data packaged in a predefined size, and kept small for easy transportation. So when large amounts of data need to be sent from point A to point B the data is broken up into numbered packets and sent to a destination. The packet may take any route to its destination, and once its arrives the data is reassembled based on its number. If a packet does not arrive at the destination, the destination only requests the retransmission of that one packet, not all the packets (TCP only).[6] All Internet communications, such as email, web browsing, and file downloads all use packets. A packet is made up of a source IP address and port (who sent it), a destination IP address and port (where's it going), the data, protocol information (rules to follow), and error checking (verification that the data arrived in the correct format).

In packet filtering, only the protocol and the address information for each packet is checked. All data within the packet is ignored. Therefore, packet filtering policies for firewalls can only allow and disallow packets based on the source IP address and port, the protocol and the destination IP address and port.

Overall, since packet filter firewalls can determine the source, destination, and the protocol being used they can effectively monitor and manage traffic coming and going from a network. However, since they do not care about the actual data being delivered this can be exploited by hackers and used to compromise systems. This type of firewall is also very susceptible to packet spoofing, which is a tactic that hackers use to change their source IP address to match an internal address, thus tricking the firewall into thinking they belong within the network.

Circuit- Level Gateways

A circuit-level gateway is a firewall that validates connections before it allows the traffic through to its destination. Unlike packet filters that allow or disallow packets through the firewall, a circuit-level gateway also determines whether the connection is valid between the source and destination, according to predefined rules before allowing the connection to proceed. Whether a connection is valid or not can be determined by the following:

- Source IP address and port
- Destination IP address and port
- Time
- Protocol
- User
- Password

Since circuit-level gateways can validate connections based on the above categories they do improve considerably on packet filtering by allowing the administrator to control the traffic with more options such time, user, and user password. By utilizing these additional control features data using the UDP (user datagram protocol) protocol becomes much more secure since the source IP address can be verified, which is not done typically for UDP traffic. Additionally, this makes packet spoofing much more difficult compared to a packet filtering firewall.[7]

Application Gateways

An application gateway is another name for a proxy, and a proxy is a device that performs data exchanges with remote systems on behalf of other systems protected behind a firewall. This is very useful since the machine protected behind the firewall is invisible to the remote machine and the rest of the world. No client ever directly communicates with the machine behind the firewall, and they never realize it. Proxies also have rules for traffic such as permitting some commands to servers but not others, limiting file access to certain types, and varying rules according to authenticated users. They also have very detailed logs of all traffic, monitored events, and have advanced features that allow alarms, and paging notifications to users should a predefined condition be met.

Overall, application gateways are the most secure type of firewall, but they are also one of the most complicated to setup and manage.[7]

State-full Inspection

State-full inspection is a firewall technology developed by Check Point Software that actively tracks all TCP and UDP connections during their life. Packets coming from the outside must have a valid connection being monitored by the firewall in order for the traffic to pass through the firewall. How does it do this? Like a packet filter, state-full inspection inspects the packet, but unlike packet filtering it also checks the data to ensure the data is what it claims to be. For example, if a user makes a request for a website, the firewall tracks that connection and when the website information is sent back to the requesting user the data is verified and then the traffic is allowed through. All other traffic that does not originate from the inside network and is not being tracked by the firewall is dropped.[8]

3.5 Hardware vs. Software Firewalls

First off all firewalls run on hardware, and use software. The difference in the term hardware firewall and software firewall is whether the firewall is a dedicated hardware appliance, meaning a hardware device with its own built in operating system and firewall software, or a server with a Windows, Unix or Linux operating system and a firewall product installed on it after.

Software firewalls can be very challenging to setup, and you must always ensure that the operating system that it's running on is fully patched or your firewall could become a security hole itself. On the other hand since the software is running on a server you can easily modify the hardware to increase performance, by adding additional processors or memory. Additionally, as your organization changes so can your firewall. Adding additional network cards and reconfiguring the rule base may allow new subnets, or rezoning of a subnet for advanced security. The cost associated with software firewalls ranges from free to thousands of dollars, depending on the manufacturer and the various options selected.

Hardware firewalls usually come preconfigured from the manufacturer and the installation can vary depending on the organization and the desired security rules. Security holes in the software are typically fixed using one patch unlike Windows operating systems for example, where it can take many patches to fix a hole. However, any changes to an organization's security policy resulting in hardware changes will cause a problem since they cannot modify the hardware. The cost associated with a hardware firewall solution ranges from a few thousand to thousands, and sometimes hundreds of thousands of dollars. It all depends on the organization, and its needs.[9]

There always seems to be a debate over which type of firewall solution should be used, hardware or software? The simple answer is it depends. There are advantages to both, so the best answer is to match up your organization with both the advantages of either a hardware or software solution and the desired budget for the project.[10]

4. Intrusion Detection

4.1 What is Intrusion Detection?

In a perfect world all our computer systems from workstations and servers to networking equipment would be 100% secure and there would be no possibility of these systems ever being attacked by a hacker. Unfortunately, we do not live in a perfect world, we live in the real world where systems are attacked daily, and operating systems constantly need to be patched to correct some new security flaw. This constant battle between security flaws and the patching of operating systems would be better if we could catch the hackers who maliciously attack systems. Thus, intrusion detection was born. The goal of intrusion detection is to detect the attack, and hopefully gather enough information that the attacker can be identified and caught. J.P Anderson, introduced the concept of intrusion detection in his technical report Computer Security Threat Monitoring and Surveillance (1980), where he defined an intrusion attempt or a threat to be the potential possibility of a deliberate unauthorized attempt to access information, manipulate information, or render a system unreliable or unusable.

4.2 Types of Intrusions

Intrusions can be divided into six main types:

1.) Attempted break-ins.
2.) Masquerade attacks.
3.) Penetration of security control systems.
4.) Leakage.
5.) Denial of Service
6.) Malicious use.[11]

4.3 How does Intrusion Detection Systems (IDS) work?

When an attempt to attack a system is made, having those attacks detected, monitored, and logged is imperative. This is exactly what in Intrusion Detection System (IDS) does. IDS is a completely reactive, not proactive system. Its purpose is to only detect the attack and capture as much information as possible. Without an IDS system it is still possible to detect attacks, but your main source of information is system audit logs. Unfortunately, searching through these logs

can be very time consuming, and is often compared to looking for a needle in a haystack.

There are two techniques for intrusion detection. The first one called anomaly detection is simply the comparison of data against an established 'normal activity baseline'. Anything that deviates from the baseline is assumed to be an intrusion attempt. The problem with this approach is false positives and the problems that result from them. A baseline is never going to be perfect, and a baseline will change over time, thus false positives occur. However, the more dangerous problem is the assumption that since false positives occur, that the intrusion is in fact a false positive when really it is an intrusion. The only way to avoid these two problems is to set thresholds for your baseline, where some deviation is allowed, while still maintaining enough integrity for the system to remain effective. [11]

The second technique for intrusion detection is called misuse detection, and is comparable to anti-virus programs where you have a program that receives regular signature files to update itself for new intrusion attacks. The main problem with this approach is the same problem that anti-virus programs face. These types of programs are only as good as their signature file, and typically a new signature file only comes out when the product's manufacturer is notified of a new attack method. [11]

So once an attack can be detected, what's the next step? Prevention. How can these attacks be prevented? Adhering to a proper security policy. These types of attacks can be avoided by implementing a proper security policy that identifies the risks, and then having the appropriate solutions set up to alleviate the identified risks.

5. Patch Management

5.1 What is Patch Management?

First off, a patch is a bug fix, security fix, or an upgrade to an operating system or application. Patch management is the process of identifying and deploying patches to various technical devices, ranging from PDA's and workstations to servers.

Computer systems need to be patched to eliminate security vulnerabilities in operating systems or software. These vulnerabilities are often identified and then a corresponding fix or patch is developed to fix the problem. Patching would not be a big concern if there was only one or two patches a year, but over the last year the number of patches has grown considerably making it very difficult to manage. Larger organizations that have not standardized on one operating system are more heavily affected by this, since they could have thousands of machines all needing different patches.

5.2 Patch Management Best Practices

Rutrell Yasin does an excellent job outlining patch management best practices in his December 2003 article "Patch management best practices" as:

- Get senior executive support.
- Establish standardized patch management policies, procedures and tools.
- Clearly assign responsibilities and provide dedicated resources.
- Create and maintain an inventory of all hardware, software, services and other technologies in use throughout the agency.
- Identify vulnerabilities and appropriate patches.
- Conduct risk assessments.
- Test each patch.
- Distribute patches effectively.
- Continue to monitor the network for vulnerabilities.[12]

5.3 Patch Management Process

For the deployment of patches to be successful a proper process is needed and must be strictly followed. Not having a process or failing to use a process introduces the potential for things to go wrong. For example, if the proper testing of a patch is not completed, it would not be possible to know that a new patch may conflict with a specialty software package the organization uses. Unfortunately, in patch management if something does go wrong it can often become a disaster very quickly. For example, the above conflict between a patch and specialty software may seem minor, but if the specialty software is on 1000 machines, and the patch was rolled out to all of those machines there could be 1000 workers unable to work. Another good example is applying a patch to a server. It just takes one bad patch to cause a server to crash and become non-functional which thus affects all users who utilize resources from that server.

As with any policy or procedure there are going to be differences due to the organization, how the organization is structured, and how the organization wants to rollout patches, but the main thing is that a process exists, it's followed and the proper approval, testing, and roll back steps are included. Below is a sample patch management process that lays out the process and the above-mentioned necessities (Figure 1).

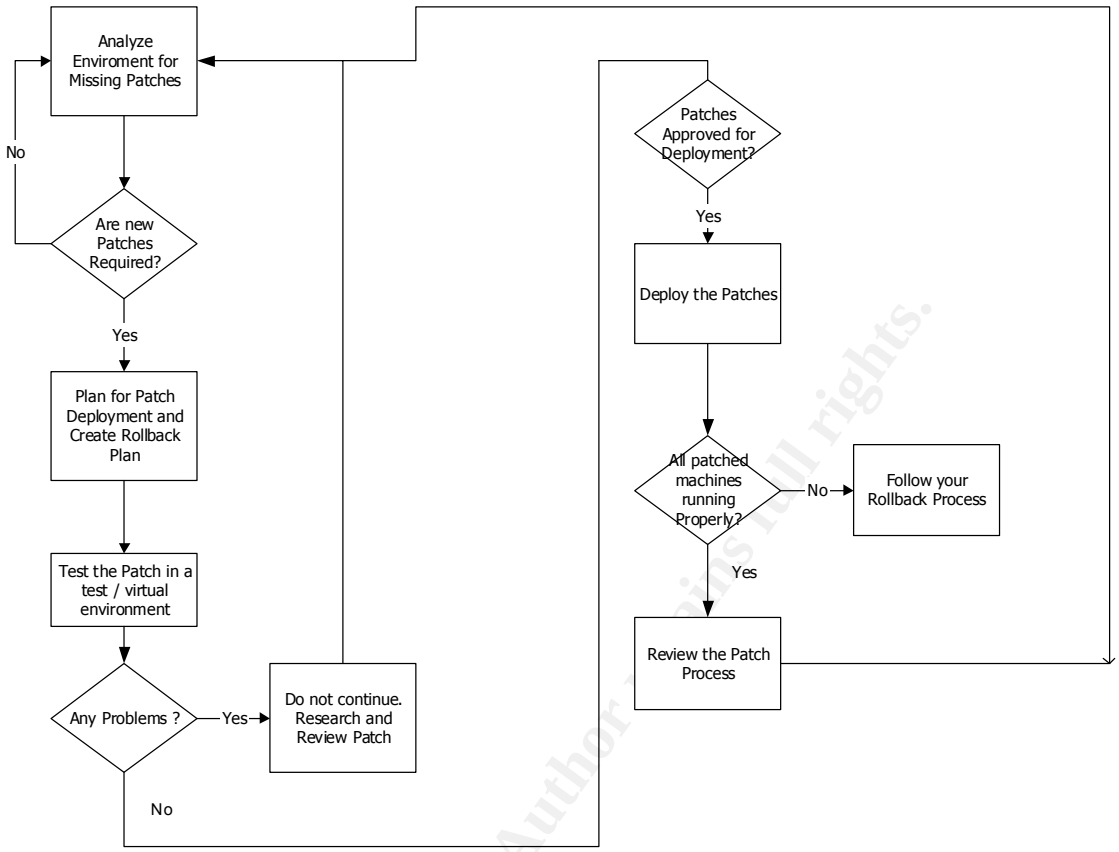

5.4 What about Remote Users?

When an evaluation of an organization is done for vulnerabilities it is important to consider the implications that remote users cause in regard to patching. When remote users connect to an organization through a VPN or a dial-up telephone pool they effectively become another machine on the network, therefore their workstations security vulnerabilities now become your organization's vulnerabilities. This is often a technical dilemma that organizations struggle with since users need to work from home, while the organization itself does not have a technical solution to protect itself. Again, this is where security policies come into place, and where a proper policy would layout the risk of having remote users connect to the network, and how this will be accomplished. This type of policy could include everything from how to connect, to what steps must be taken before your allowed to connect, such as ensuring the remote workstation is patched and has anti-virus software installed.[13]

6. Anti-Virus Management

6.1 What is a Virus and What is Anti-virus Management

A computer virus is defined by Princeton's Wordnet 2.0 Search, as a software program capable of reproducing itself and usually capable of causing great harm to files or other programs on the same computer, "a true virus cannot spread to another computer without human assistance". Up until recently this definition was accurate, but just as the whole computer industry has been developing and growing, so has computer viruses. Over the last year, computer viruses have become more intelligent and the need for human assistance is no longer needed. A good example of this was the MS Blaster virus that infected millions of machines across the world, launching a denial of service attack against Microsoft. Once the virus was in an organization's network, all the machines if not properly patched or virus protected were infected within hours, and any new machines added to the network were infected within minutes. This type of attack changed not only how viruses launched, but also served as a wake up call to organizations that Anti-Virus Management was a mission critical component of network security and therefore should be treated accordingly.

Anti-virus management is more than simply ensuring that your organization has virus protection software on the workstations. It's about a complete plan for the organization that encompasses workstations, PDAs, servers, and the organizations firewall. In addition, anti-virus management also leverages from various security policies how an enterprise will protect itself with its chosen anti-virus software, which may include real time protection, manual virus scans and how an organization will handle new virus definitions.

6.2 What Should be Protected?

All devices that have an operating system can be potentially attacked if a security flaw is discovered. Therefore, anti-virus software should be installed on all devices capable of running the software. As mentioned above, this would include workstations, PDAs, and servers. An organizations 'chokepoint' firewall is a perfect location to have anti-virus protection since all traffic must pass through this point. Comparing the traffic with known virus patterns at this point can significantly reduce the number of viruses that make it into a network. Having virus protection at both the entry location on the network and the various devices behind the firewall provides a layered approach that increases protection. Should one layer fail, there is still another layer of protection.

As mentioned with regards to patch management, remote users accessing an organization's network are a dilemma. No matter which way they connect to the organizations network if their anti-virus software is out of date, they are now a security threat on your network. A proper policy requiring these remote users to have anti-virus software is necessary.

6.3 Anti Virus Options and Concerns

Ensuring that viruses do not penetrate a network takes more than simply anti-virus software. It takes an effective plan to utilize the software and its various features and the proper policies to ensure the plan is utilized. An effective plan would lay out where and which devices should have the software installed, how the software should be configured, who is responsible for ensuring the plan is being carried out properly, and what should be done when a virus does infiltrate the organization's network.

One major issue that does take some additional planning is new virus definitions. New virus definitions typically do not cause problems, but they can and therefore should be tested prior to a rollout. Unfortunately, the dilemma is time. When a new virus hits, organizations typically do not have a lot of time to test before a virus may hit them so it typically comes down to a quick risk assessment. Additionally, if possible a staged rollout is probably best since it will allow an organization to check for possible problems if little to no testing was completed, and it will also not overload the network by trying to rollout an update to to many users at once. Typically this only applies to vary large organizations. However, since new virus definitions can be updated 3 or 4 times daily when a major virus is loose, having all the machines update themselves 3 or 4 times could cause some network congestion.

7. Physical Security

7.1 Overview

Physical security in the past was often overlooked or forgotten about by organizations until the terrorist attacks of 9/11. Organizations would spend thousands of dollars on firewalls, and other security devices but would spend very little in protecting their physical assets.

To start the process of evaluating an organization's physical security, it is important to review the security policy, the identified processes and assets that are crucial to the organization and then identify what they're physically dependent on. This can range from environmental concerns such as power availability, fire hazards, temperature regulation, and water damage to the human factor of proper passwords, locks, security cameras, and human mistake.

Once a list of all the identified dependencies has been created think of anything and everything that could go wrong with them. Then try to put the proper safeguards in place to ensure that these possible disasters never occur.

Through this process it is also important to ensure that a proper contingency plan is developed, tested and kept current for all the critical processes and assets. A

disaster can occur at any time, but being prepared and having a thorough plan in place is the difference between a surviving organization and a dead one.

7.2 Environment Issues

When starting to look at environmental issues, organizations should start looking at these issues in steps. Step one, location. Does the location of your organization provide the desired security level, or is there an environmental concern. A good example of this is running a major computer company in Florida. Does it make sense to have all your computer operations equipment located in Florida with the number of hurricanes that are ravaging that side of the coast? It is also important to choose the building location within the city. Are areas of the city prone to power blackouts, theft, fires, vandalism? The second step is the building itself. Does the building meet the organizations needs? Are there adequate power generators to help regulate power? Does the building have a UPS system should a loss of power occur? What does the building use for fire suppression, and what about water damage? Buildings can have old pipes burst, and what if a fire sprinkler goes off? These are all important components to choosing a good location and protecting your assets.[14]

This does not mean that an organization has to have a new facility that meets all of the above needs, because realistically that's not typically possible. The main thing is to identify the problems, and develop a contingency plan that explores each of them and what needs to be done should one of them occur.

7.3 The 'Human' Factor

With the human factor its important to start with the perimeter of what the organization is trying to protect. This could be the building, or could be a server room. It all depends on the business of the organization and how secure they want to be. In each case you would start with proper procedures for entry. This could be as simple as picture identification, or reporting to a security guard and as complicated as state of the art biotech thumb scanners. The type of security should match with what the organization is trying to protect. An organization would obviously spend more than a thousand dollars on protecting a million dollars worth of equipment. Once an individual (who does not work for the organization) is past the perimeter, security should still exist in the form of an escort, or video monitoring. If we are talking about a building, maybe video surveillance is adequate, but if were talking about a server room, both video surveillance and an escort is required.

With regards to server rooms, it's important to make sure the room itself is completely secure. This ranges from making sure there's no other way in, such as crawling under floors or through air ducting to making sure the above mentioned environmental issues have all been identified and secured.[15]

When discussing the human factor it's also necessary to discuss human error. Everyone makes errors, and accidents do occur such as pulling out the wrong network cable or tipping over a loose power cable. How do you stop these accidents or errors? You can't. You can try to avoid human errors while working on important systems, by following processes or procedures or working with partners, but most times there impossible to stop.

7.4 Technological Safe Guards

Since accidents, errors, or problems with computer systems and operating systems can occur any time, technological safe guards such as making services redundant should be considered for any critical systems. For example, a company that provides web hosting would not have only one server providing all its web hosting, they would have many servers, and therefore can suffer the loss of one or two servers. Where possible, make services redundant.

7.5 Contingency Plans

All of the above potential problems, and solutions should be properly documented in case a disaster occurs. The main premise behind this plan is to describe all the steps needed to get the organization back up and running. This would include the right contacts, alternate facilitates, where server, network and data backups are kept.
This type of plan should be developed with participation from across the organization, and will be a living document always changing as the organization continues to change.

8. Conclusion

It's important to note that all of the six categories discussed above are directly connected in what's termed 'defense in depth'. This concept is that security is layered, and even if one layer has been compromised additional layers of security exist. For example, a properly configured firewall should stop the majority of new viruses from entering a network, however if a virus does make it past the firewall (which will happen), patch management and anti-virus management are there at the next layer to hopefully contain and stop the virus from spreading.

Throughout this paper, the various layers of 'defense in depth' were introduced and discussed to varying degrees. The goal was to provide a new network administrator or security officer a basic introduction to the various topics. The topic and the complexity involved determined the level of detail provided in this paper. In no way should more detail about one topic, and less about another reflect a level of importance. Each topic represents a vital component of network security, and is required in a complete network security architecture.

In conclusion, some type of network security is usually present in an organization. Unfortunately, organizations often rely too heavily on one technology, such as a firewall for their complete defense when they should start at the beginning, determining what needs to be protected. As it was shown throughout this paper, once the basic question of what to protect has been answered it's often easier to determine how to protect it and how much protection is needed. Instead of hoping for the best, that a firewall for example will stop viruses from entering a network, plan for it. Plan for what is not expected to happen. Utilize multiple layers of security to safeguard a network. Remember, network security is imperative in the world of today. Failure to properly secure a network, or an organization can seriously affect a businesses ability to operate or even survive in today's technological age.

9. References

1 Kok, Kee, Chaiw. "Security Policy Roadmap – Process for Creating Security Policies." SANS Institute. 2001.
URL: http://www.sans.org/rr/papers/index.php?id=494 (October 8, 2004)

2 Wills, Laura. "Security Policies: Where to Begin." SANS Institute. 2003.
URL: http://www.sans.org/rr/papers/index.php?id=919 (October 8, 2004)

3 Chapman, Brent D. and Zwicky, Elizabeth D. Building Internet Firewalls. O'Reilly, 1995. 48.

4 Cafarchio, Pete. "Firewalls: Are We Asking Too Much?" Computer Security Institute. 1998. URL: http://www.spirit.com/CSI/Papers/fw-ask2much.html (October 8, 2004)

5 Kunene, Glen. "Perimeter Security Ain't What It Used to BE, Experts Say." Jupitermedia Corporation. (2004).
URL: http://www.devx.com/security/article/20472/1954?pf=true (October 8, 2004).

6 Sheldon, Tom. "Firewall." Big Sur Multimedia. 2001.
URL: http://www.linktionary.com/f/firewall.html (October 8, 2004).

7 Avolio and Blask. "Application Gateways and Stateful Inspection." Trusted Information System, Inc. 1998. URL: http://www.avolio.com/papers/apgw+spf.html (October 8, 2004).

8 Webopedia. "What is stateful inspection? – A Word Definition From the Webopedia Dictionary." URL: http://www.webopedia.com/TERM/S/stateful_inspection.html (October 8, 2004).

9 Shinder, Deb. " Comparing Firewall Features." Windows Security. 2004.
URL: http://www.windowsecurity.com/articles/Comparing_Firewall_Features.html (October 8, 2004).

10 Pacchiano, Ronald. "Firewall Debate: Hardware vs. Software." Small Business Computing. 2003.
URL: http://www.smallbusinesscomputing.com/webmaster/article.php/3103431 (October 8, 2004).

11 Sundaram, Aurobindo. "An Introduction to Intrusion Detection." Association for Computing Machinery. 1996.
URL: http://www.acm.org/crossroads/xrds2-4/intrus.html (October 8, 2004).

12 Yasin, Rutrell. "Patch management best practices" Federal Computer Week Dec. 01, 2003. URL: http://www.fcw.com/fcw/articles/2003/1201/cov-patch2-12-01-03.asp (October 8, 2004).

13 Gadue, David J. "Using Proactive Depth in Defense to Ease in Patch Management Problems." SANS Institute. 2004. URL: http://www.sans.org/rr/papers/8/1449.pdf (October 8, 2004).

14 Fickes, Michael. "Bridging the Gap." Access Control & Security Systems. 2004.
URL: http://securitysolutions.com/mag/security_bridging_gap_2/ (October 8, 2004).

15 University of Chicago. “NSC: Physical Security.” 2000.
URL: http://security.uchicago.edu/docs/physicalsec.shtml (October 8, 2004).