Threat – Risk Analysis

For:
SEC703
Advanced Security

By:
Megan E. Cutler
Jaanus Anja
Boris Plotkin
and
Daniel Sowinski

December 6, 2009

Table of Contents

Purpose

The purpose of this document is to illustrate the risks posed to an IT environment and highlight why it is important to be vigilant and maintain a high level of security in any such environment. On the surface, the consequences of negligence in securing the IT environment may seem insignificant. However, with a more in depth look it is easy to see that security is a critical component to any IT environment, most especially in any business setting. In society today, every company relies on computer systems to store, file, transmit, and share information, whether it be trivial communications or sensitive business critical information. Considering the deep trust we put into our computer systems to provide us ease of access to critical services, it is important to ensure that all such transactions are undertaken in a secure, trusted environment which will not open the company, its clients, or its customers to to the risk of theft or other such cyber attacks. In this document we will highlight the most critical security concerns related to the company's IT environment. It is important not only to understand the risk factors involved, but to understand how to protect the company and its customers from these and other technological threats.

Executive Summary

The first thing this document focuses on is an analysis of the risks and threats posed to the company by its current IT environment. Each area of the IT environment is broken down so that a high level of focus can be given to the risks and threats posed to each one. We will assign a risk level to each factor examined in order to illustrate how valuable the resource is to the company and how much damage it would do to the company if it were to be compromised. We will be focusing on such areas as the servers which are used by the company, the services which are offered by each server, and the clients which are used to access them. A smaller level of focus will be given to the network itself, however, where the Internet is involved, network security should always be considered a factor.

When we have finished analyzing the risks posed to the company's IT environment, we will examine ways in which these threats can be prevented. It is important to offer a secure environment for all IT transactions to take place in and we will look at some simple an effective ways to turn an insecure environment into a place where the company can be assured of security. Though we will not focus on specific tools, the principals we will highlight can be applied to the current environment and any new updates or additions to ensure that security is always maintained at a high level.

Though we will not be focusing a great deal on specific security products, we will demonstrate how to configure and deploy some resources which are valuable in protecting an IT environment, such as an IDS (intrusion detection system) and a Firewall. Intrusion Detection Systems can be difficult to set up initially, but the value they provide is well worth the effort. We will provide a guide to getting started, to ensure that the network snapshot we have provided can be achieved with minimal effort.

Finally we will discuss how to maintain a secure environment, once all objectives have been achieved. Security systems can never be left alone for very long before they become useless. To ensure a high level of security, it is important to maintain security systems, provide updates for new security releases and exploits, and ensure that you are always one step ahead of the attacker. We will take a brief look into considerations for the continued safety of the company's security environment.

Threat – Risk Analysis

We will begin by taking a look at the risks which threaten the company's current IT environment. The environment will be broken down into pieces in order to be more easily examined. The factors which will be considered include: company servers and the services which they offer, client computers which access company resources, and the network which is used to connect all parts of the company together. In order to properly understand the risks involved, it is important to first establish a baseline. What are risks and what are threats?

Threats are anything which poses a danger to the system or network. In this case a threat is anything which could cause the company's business or productivity to be interrupted. *Vulnerabilities* are weak points in the system. Vulnerabilities are what allow threats to become dangerous. An *exploit* is anything which takes advantage of a vulnerability. Exploits are what will cause real damage to the company's systems if they are allowed to be executed. *Risk* is the likelihood that a vulnerability will fall prey to an exploit. The risk posed to the company by a vulnerability depends on many factors including but not limited to; the amount of damage the exploit will cause to the vulnerable system, the amount of effort required by the attacker to take advantage of the vulnerability, the direct payoff to the attacker for the amount of effort put into the attack, and the amount of money it will cost the company to recover from an exploitation of the vulnerability. The higher the damage to the company, the higher the risk level is said to be.

Windows Servers

We will begin by taking a look at the company's servers since they serve a critical role in the company by providing access to important company resources such as internal databases, web sites and email. Currently the company has two servers; a Windows server and a Linux server. Each system provides its own individual security risks, so we will look at them individually. Currently the Windows server runs several services such as IIS (Internet Information Services; a Microsoft Web Server), a Windows Mail Server (SMTP) and an FTP (File Transfer Protocol) server. These services provide access to the company website for both internal and external users, email services for employees, and the ability for employees to share information with each other. Many of these services are critical functions; if they were interrupted there would be a noticeable impact to productivity for the company. We will begin by taking a look at the risks to the current server configuration and the resulting risks to the company from the same server.

Threats to the Server

Servers represent centralized collection of tools, applications and services critical to the administration, maintenance, and day to day running of an organization's network infrastructure. Service disruption of any server providing mission critical services can potentially affect every user on the company network. An outage to this server would mean the disruption of email services to all employees whose email is managed by this server. Likewise, an outage of the web page could cause the company to loose business transactions from frustrated customers. The loss of money and productivity represents a significant loss to the company.

The server is also at risk to viruses, malware, trojans and other exploits which target the Windows operating system. While every operating system is vulnerable to exploitation by viruses, Windows is most often the target of worms. Viruses can pose a wide range of threats to any system they infect; they can cause damage to the server itself causing a loss of data and a great deal of time to repair and reconfigure the system. Services can be 'hijacked' by an attacker and used for nefarious purposes.

Aside from the server as a whole, the individual services hosted by the server must also be considered. For instance, there are threats to the company's web server hosted on the Windows server. If the company website becomes inaccessible to customers they may take their business elsewhere, causing the company to lose money. New customers may not be able to discover the value of the services offered because the website is unavailable during their search, again causing the company to loose business. Web resources which employees depend on may become inaccessible causing a loss of productivity for the company as well. If an attacker were successfully able to compromise the web server, rather than make it unavailable, they could upload malicious code which may cause those who visit the site to be infected by a virus. An attacker could even cause the site to redirect to another malicious page or to other undesirable web page. Any such disruption could cause the company's reputation to suffer in the eye of its customers and they may become distrustful of accessing the company's online resources in the future.

Likewise the company's mail server, also hosted on the same server, would cause an interruption in communication between employees and between the company and its customers if it were to become unavailable. Time critical communications would be unable to arrive at their destinations in time. Transactions could be put in jeopardy if the issue were not able to be solved in a timely manner. Furthermore, if an attacker were able to steal sensitive information from the email server, the company could be liable for the loss of that information, most especially if the information stolen is critical business information or private information being provided by a customer.

Threats from the Server

Once a server has been compromised, the attacker may be able to use that server as a launching point to do more damage to other systems. They may decide to turn their eyes inward to attack other systems within the company or they may decide to turn their malicious intent onto other external systems. For instance, if the company's mail server were to be 'hijacked', the attacker could use it to send spam messages. Like the compromise of the web server, such an attack would cause harm to the reputation of the company.

Another security risk posed by the current configuration is the inclusion of an FTP server along with an email and web server. While it may seem logical since FTP also requires access to the Internet, FTP is known to open several security vulnerabilities to the system. Often FTP is misconfigured or allows access to anonymous users if it is not properly locked down. Since FTP is much easier for an attacker to gain access to than many other services, it is most likely this service they would chose to attack. Once they have gained access to the system via the FTP server, it would be much easier for them to access other systems hosted on the same server.

When considered separately, each service hosted on the server has its own risk level:

- The Web Server has a Risk Level of **Medium** – interruption to the service may cause a loss to the company, but it is less likely to be business critical than other functions such as email.
- The Email Server has a Risk Level of **High** – interruption to the normal flow of communications could cause loss of business for the company and bring productivity to a stand still. Email is a critical service for communication between employees and between the company and customers and therefore should be protected as much as possible from service outages.
- The FTP server has a Risk Level of **Low** – interruption to the FTP service is less likely to be a critical issue. Since the FTP server has much lower security considerations than the other two services offered by the server, it should not be included on the same server as either.

When all components of the server and the services it offers are considered together the server has its own risk level. The risk level of the server is **High** because it contains business critical services which would cause damage and cost the company money were the server to be out of commission for a long period of time.

Linux Servers

A second server, this time on a Linux operating system, hosts a second set of services which are no less central to the core operations of the company. Though there are some different security considerations based on the server's operating system and the services it is offering, it is important to note that many of the same vulnerabilities will apply to this server as well. Currently the Linux server offers DNS (Domain Name Service) services to all employees which allows them to access the Internet on a daily basis. The company database is also hosted on this server using an SQL (Structured Query Language) server. It also provides TFTP (Trivial File Transfer Protocol) services and is host to VNC (Virtual Network Computing) which allows remote viewing of the server system for administration purposes. Like the other server, we will take a look at the risks posed to this system, and the risks which it poses to the company.

Threats to the Server

Though Linux operating systems tend to be targeted by viruses less often than Windows systems, they are not immune to viruses, malware and other such vulnerabilities. They too can be infected and damaged by such systems, causing downtime and a great deal of effort to recover and reconfigure if the virus results in the loss of data, or damages the system. Like damage or outage to the Windows server, interruption of the server's services can cause significant loss in productivity from employees, reduction of sales from customers, and damage to the company's overall reputation.

The DNS service provided by the Linux server is required in order for employees to access the Internet. Interruption of the server or of this particular service will cause all employees to loose access to the Internet and any web resources they require to do their work. This could cause employees to be unavailable to customers or clients. It is also highly likely to result in a high amount of lost productivity, as many employees will be unable to do their jobs without access to the network or the Internet.

Since the company database is hosted on this server, it is critical that employees which require access have access to the information in the database at all times. If the database is unavailable orders will not be able to be processed, invoices will not be able to be issued, and employees will not be able to receive their pay. Like the email server, the company databases houses sensitive information which the outside world should never be able to access. Loss of this information poses a great deal of problems for the company. Theft of critical company information or private information gathered from customers will leave the company liable under legal regulations, not only causing harm to the company's reputation but likely costing them a great deal of money.

Threats from the Server

As with the other server, once it has been compromised, the attacker can then use the server as a launching point for other malicious activities. Since this server houses the DNS services for the company, the attacker may be able to use a method known as “DNS Cache Poisoning” to spread further malicious content. Cache Poisoning allows the attacker to redirect requests for legitimate websites to any website of their choosing. It may be a website which looks exactly like the real website but has malicious intent. It may be a website which houses a virus, or it may redirect to other undesirable locations.

If the company database is open to access from the Internet and is not properly protected, it may fall risk to SQL Injection. This form of attack allows a hacker to determine information from the database by inserting their own queries into applications which perform queries against the company's database. SQL injection can also result in damage to the database if the attacker inserts their own information into the database. SQL injection techniques can also be used to break into the server by inserting malicious code which is then executed on the server to interrupt the program or allow the attacker to gain access to the system. As with the other server, once an attacker gains access to one service running on the server, they will be able to use that position to gain access to everything else hosted on the same server.

Like FTP, the TFTP service is insecure and introduces risks to the system which should be avoided due to the sensitive nature of the material stored on the server. The same can be said of VNC. The VNC service is known to be prone to exploitation and it is not recommended to install VNC on servers which host sensitive information.

When considered separately, each service hosted on the server has its own risk level:

- The DNS server has a Risk Level of **High** – if the company were to be cut off from the DNS service the inability to reach the Internet would cause a high loss of productivity and could cause damage to company profits depending on the length of the outage.
- The SQL server has a Risk Level of **High** – if the data stored in the internal database were to be compromised, the company may be liable for stolen customer or client information and may suffer further losses if business practices are exposed to competitors.
- The TFTP server has a Risk Level of **Low** – Interruption to the TFTP service is likely to go unnoticed
- The VNC server has a Risk Level of **Low** – Loss of the VNC program is unlikely to cause a large impact for a large number of users.

When all components of the server and the services it offers are considered together the server has its own risk level. The risk level of the server is **High** because it contains business critical services which would cause damage and cost the company money were the server to be out of commission for a long period of time.

Client Computers

While the individual computers used by company employees may seem insignificant when compared with the critical servers that make up the core of the company, they cannot be ignored where security is concerned. Each user depends on his or her computer in order to do their work. Whether they are replying to customer emails, scheduling important meetings, or preparing presentations for important client meetings, most employees will find themselves relying on their computers to ensure their work is done in a timely manner. Because the computers employees use are also connected to the company's network it is important to realize that threats can be transmitted from an individual employee's computer to any other computer on the network (including the servers), or can pose a threat to the network as a whole.

Threats to Clients

Since each employee relies on their computer to do their job, an interruption to the use of their computer represents a loss of productivity to that user. However, since the computer is only needed by one person, an interruption to the system is relatively small in scale; it will only affect one user instead of many. For this reason it is not as big an issue as a server outage, but the interruption of many client computers can quickly add up to a large number of unproductive workers. This is why it is important to consider security for client workstations as well as company servers.

Like their counterparts, user's computers are susceptible to viruses. In fact employee computers are more likely to be infected with viruses because they will more often be used to access the Internet. It is inevitable that an employee's computer will be used to browse the Internet for personal use at one time or another. This can inevitably lead to malware or viruses making their way onto the computer. Some viruses are easy to catch as they are flashy and make themselves known through pop ups. Others hide in the computer's system files and the user is never aware of them until irreparable damage is already done.

Threats from Clients

While the employee's computers will not be running services which are used by the entire network, they will be using applications which the user will rely on in order to complete certain tasks. Every application installed on a computer can represent a security risk. Since applications are all maintained by different companies and because some may offer patches less frequently than others, or users may even neglect to update new releases for the software, it is important to be aware of the security risk that installation of applications represents. Users are likely to install software which they want to use for personal reasons rather than business reasons, such as Skype or iTunes. Each of these programs brings a new security risk along with it and this risk should be carefully considered before the application is installed. As with the servers, if an attacker is able to gain access to a user's computer via an application they have installed, it will not be hard for them to gain access to the rest of the system.

Once a user's computer has been infected with a virus or exploited by an attacker, the computer then becomes a risk to the rest of the company, especially if it is plugged into a secure network after the infection takes place. The virus may try to replicate itself to all other computers attached to the network. Depending on the type of virus, it may use the computer to conduct attacks against other computers or company resources. It may turn the computer into a “zombie” which is controlled by the original attacker. It may even try to copy itself to other computers, turning them into “zombies” as well. These “zombie” computers can then be used to carry out denial of service attacks against company servers or other external targets as part of what is called a “botnet”.

Though it may seem odd, users also pose a risk to the computers they use. Even the most sophisticated Intrusion Detection System cannot always protect against social engineering techniques known as “pharming” and “phishing”. Though many people do not know what “phishing” is, just about everyone has received a “phishing” email. We have all had spam messages informing us that we need to “verify our account settings before the account is deleted”. Though most companies which provide services which are used for “phishing” attacks release notifications that they will never ask for your password, many people still fall prey to these types of attacks (or of course attackers would no longer attempt them). If a user gives away account information for a sensitive internal system, the attacker may be able to use that password to gain access to internal company resources. Even something as simple as giving away the password that is used to access the user's single computer can pose a risk, as once the attacker gains access to the computer he may be able to find ways to access other systems in the company's internal network.

When all risk factors for client computers are considered, client computers can be said to have a Risk Level of **Low** – an interruption to a single workstation will only affect a single user and will not cause a wide disruption of service or productivity. It is also relatively easy to prevent viruses from becoming a major problem.

It should be noted, however, that the risk from social engineering may be considered **Medium** depending on the type of attack.

Recommendations

By now it is clear why security is such a concern for any IT environment. With so many potentially malicious threats lurking around, it is impossible to trust unknown users not to tamper with precious company resources. Without protection, any one of the scenarios above is a possibility. Servers left completely undefended, would be like a castle with a permanently opened drawbridge. System administrators would have no way to identify if an attack was taking place or if an attack had recently occurred. Most attacks come with early warning signs; attackers will attempt to port scan a target system in order to identify possible vulnerabilities. Attackers often gather as much information as possible about a target before they attempt to take control or do any real damage. Without protections, these early warning signs will go unnoticed and an attack may not be noticed until it is too late to prevent or even reverse the damage. The net result of unknown information gathering attempts in the worst case scenario is an attack against a server with a service that has a known vulnerability to exploit. An attack of this nature could be little more than cyber-vandalism (changing the background of a home page), to compromising and installing a backdoor or rootkit, programs which enable the attacker to return later and use the compromised machine as desired.

Now that we have discussed the dangers to the IT system, we will discuss ways in which such attacks can be prevented. While there are no specific programs listed here, these principals can be applied to any component of the network in order to ensure better security. The same principals should be applied to all changes in the network to ensure uniform security. As we have seen, a security risk to any one portion of the network could prove to be the downfall of any other portion of the network.

Protecting Servers

It is essential to protect the company's servers because they often lay at the core of operations. Often they host sensitive information which must be kept safe from loss or theft. While backing up data may protect against loss, it does nothing to prevent theft. Since many of the services hosted on company servers are required for day to day operations, the best thing to do is ensure that they can be kept running and down time, in the event of a problem, can be kept to a minimum. Here are some good ways to protect servers on your network:

Eliminate Weak Links:

- Disable all services which are not absolutely required. Any service which is running but not needed is another stepping stone an attacker can use to gain access to your system.
- Separate services according to security concerns. If one service has a Risk Level of High, then it should not be installed with a service whose Risk Level is Low. High risk services require as much locking down as possible and no service which is prone to vulnerabilities should be allowed to share the same server as these critical systems.
- Disable guest accounts.

- Rename or disable Administrator accounts to make them harder to find or access.
- Disable anonymous accounts for any service which allows them (such as FTP) and require users to have accounts and passwords to login.
- Avoid default installations.
- Servers hosting internal company databases should not be allowed to connect to the Internet.
- Servers with services which must reach the Internet should be protected by DMZs (Demilitarized Zones)
- No session to browse the Internet should ever be opened from a server system.
- Do not allow services to run as the administrator (or "root" in the case of Linux Servers)
- For Linux servers, run services in a "chrooted" environment where possible. "Chroot" environments prevent services from running at the highest level of authentication (root).
- Allow access to resources and systems according to the rule of least privilege

Take Preventative Action:

- Keep up to date with information about new threats which are released regularly on news groups and vendor websites
- Apply all the latest vendor patches and operating system updates
- Be discreet with sensitive information
- Use an Intrusion Detection System and regularly audit the logs to ensure that an attack is caught as quickly as possible
- Allow the Intrusion Detection System to actively modify the firewall to cut off an attack as soon as it is detected

Have a plan for response to attacks:

- Bring compromised servers offline or remove them from the network as soon as an issue is discovered
- Check all logs to determine what actions where taken against your server and document as much of the activity as you can.
- Do not change the system until an investigation can be completed.
- Restore data from backups.
- Have a written plan of action for compromises and ensure that it is followed. This document should include a clear "chain of command" of who should be contacted and consulted in the event of a breach in security.

Protecting Client Computers

As we have illustrated, there is value in protecting client computers. Client computers can often prove easier to protect than servers, and implementing these preventative measures often saves a lot of headaches in the end. Keeping client computers safe will also help to improve the security of the internal network and take some of the burden off of the network measures put in place to do so. Here are some useful ways to protect employee computers:

Grant access to resources, accounts and computers based on the rule of least privilege:

- Make sure each host only has access to network resources which are *required* by the user and no more
- Lock down user accounts so that the user cannot do things that will open the computer to security vulnerabilities. An example of this is to prevent installations of client applications known to have security vulnerabilities. This can be done via group policy

Ensure that security is considered when making changes to client computers and policies:

- All new installations and upgrades should be tested before being pushed to client machines
- Avoid installation of new programs which remain relatively untested

Eliminate Weak Links:

- Enforce strong passwords. Strong passwords are at least eight characters long, contain upper and lower case letters and include either numbers or symbols or both. Strong passwords should not be based on dictionary words and should not increment (IE you have a number at the end of your password and increase it by 1 each time you need to change your password). The more complex a password and the less like a dictionary word, the harder it is for an attacker to guess.
- Enable an account lockout threshold so that an account will lock after a set number of failed login attempts. A good number to use is five attempts; this allows for a user to get into their account if they initially may have forgotten their password (perhaps because they changed it or just returned from vacation) but does not allow enough attempts for an attacker to break into the account.
- Log attempted login failures and regularly audit the logs so that you can be aware of attempts to break into the system.
- Make passwords expire after a certain period of time. This limits the amount of time an attacker has to guess a password or crack passwords that have been stolen from a compromised machine.

Take Preventative Action:

- Educate the user on how to avoid social engineering tactics such as phishing.
- Make users aware of what applications should be avoided due to security risks
- Install antivirus software which regularly scans the system
- Keep antivirus software up to date
- Regularly install security patches and updates
- Configure a local firewall and make sure every client has it enabled
- Install and maintain a host based IDS on every client

Have a plan for response to attacks:

- Inform users of what to do if they think their computer has been compromised
- Have an avenue of communication open to users who think their computer has been compromised such as an IT helpdesk, helpline or an IT SPOC (single point of contact)

Protecting the Network

Though until this point we have only touched on the network in passing, it is as important to protect the internal network as it is to protect the machines which are attached to it. Ensuring that only legitimate network traffic is allowed to pass into the internal network will keep the systems attached to it more secure and ensure that all information sent within the business stays secure within the boundaries of the company's internal network. As we have shown before, a breach to a computer inside the network could prove to be a threat to the rest of the internal network. Here are some ways to keep your internal network secure from outside attacks:

- Use VPN connections to ensure secure connections to company resources
- Disallow any external communications to internal servers that do not come through VPNs or implement a DMZ to “speak” to external networks such as the Internet
- Where possible, encrypt network traffic
- Use a network based IDS to detect attacks and cut them off if an attack is detected. Regularly

audit the logs to be aware of attacks or attempts to probe the network

- Use network firewalls to filter and block traffic
- Make use of a switched network in order to prevent data from being intercepted by sniffers

Implementing Quality of Service

Security is not the only concern for the company network. While hardening the system against attack is critical, quality of service must also be balanced against protection. If too many security measures are implemented or implemented poorly, it could bog down the network creating bottle necks and slow downs. A secure network is no good if the employees who need it cannot make use of it. Not only will a bogged down network reduce productivity, but it would make it easy to perform a denial of service attack and bring the network to a crashing halt. Implementing quality of service measures along side network security is a good idea. Block network activity which is not related to company business or lower the priority of such traffic so that it will not interfere with business functions.

Network Diagrams

We have created a diagram to illustrate some of the security principals which we have discussed for securing the network. To properly demonstrate the difference between a secure network and an insecure network, we have first created this diagram of the current network infrastructure:

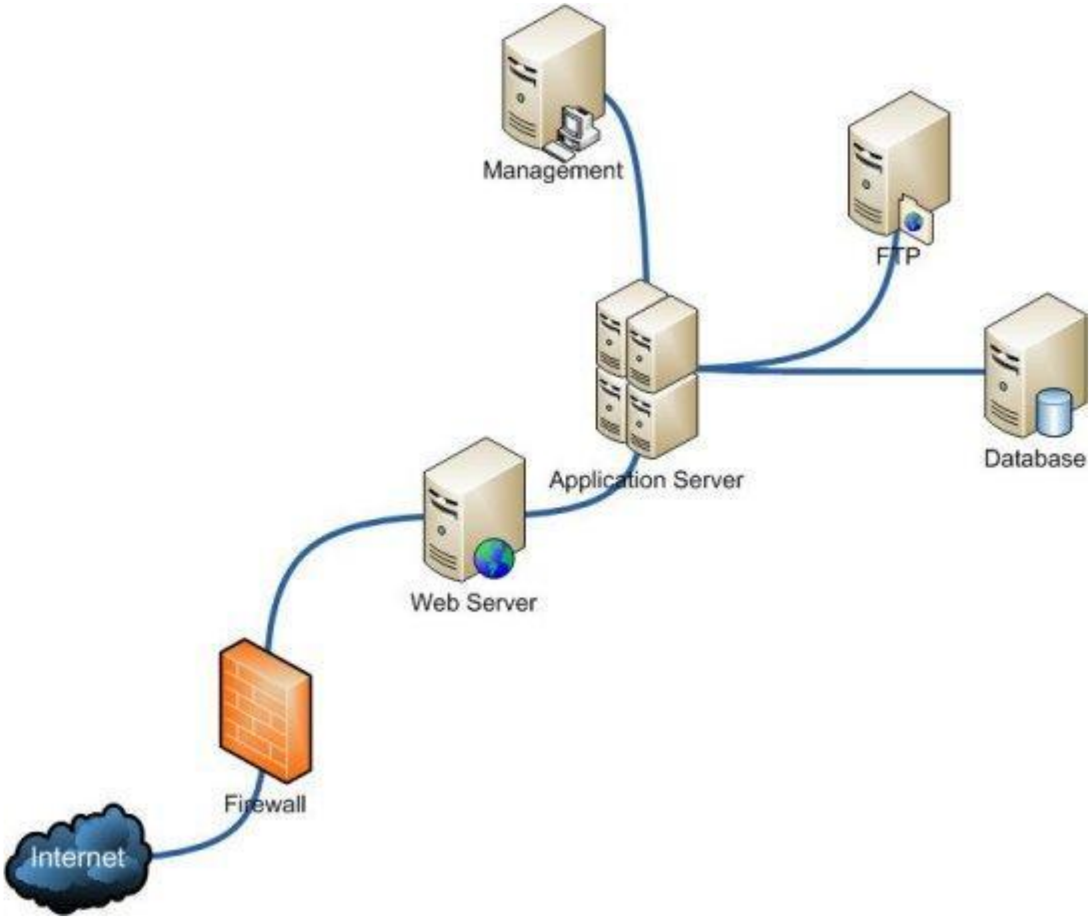

As you can see, all of the servers are connected together on a single line. Several services are hosted on a single server and there is very little protection laying between the servers and the Internet beyond. Currently, a single firewall serves to protect the network. Should that firewall fail, there would be no protection for the network at all.

The following diagram, in contrast illustrates many of the concepts we have tried to describe. Routers and switches have been implemented to route network traffic, help block malicious traffic, and ensure that an attacker cannot “listen” to network communications. An Intrusion Detection System is working to actively protect the network. If desired, as second Intrusion Detection System could be implimented in case the first one should fail. A DMZ has been implemented to sit between the internal and external networks. This ensures that the inside and outside can communicate with each other, but that the external Internet has no way of communicating with the protected internal resources directly. This also ensures that something sits between the internal server and an attack. If an attack were to occur, it would most likely target the DMZ, rather than the server which hosts the master copy of the company database. Finally, separation of services to different servers has been implemented to ensure security of sensitive information. Backups have also been applied to ensure that there is no single point of failure for servers.

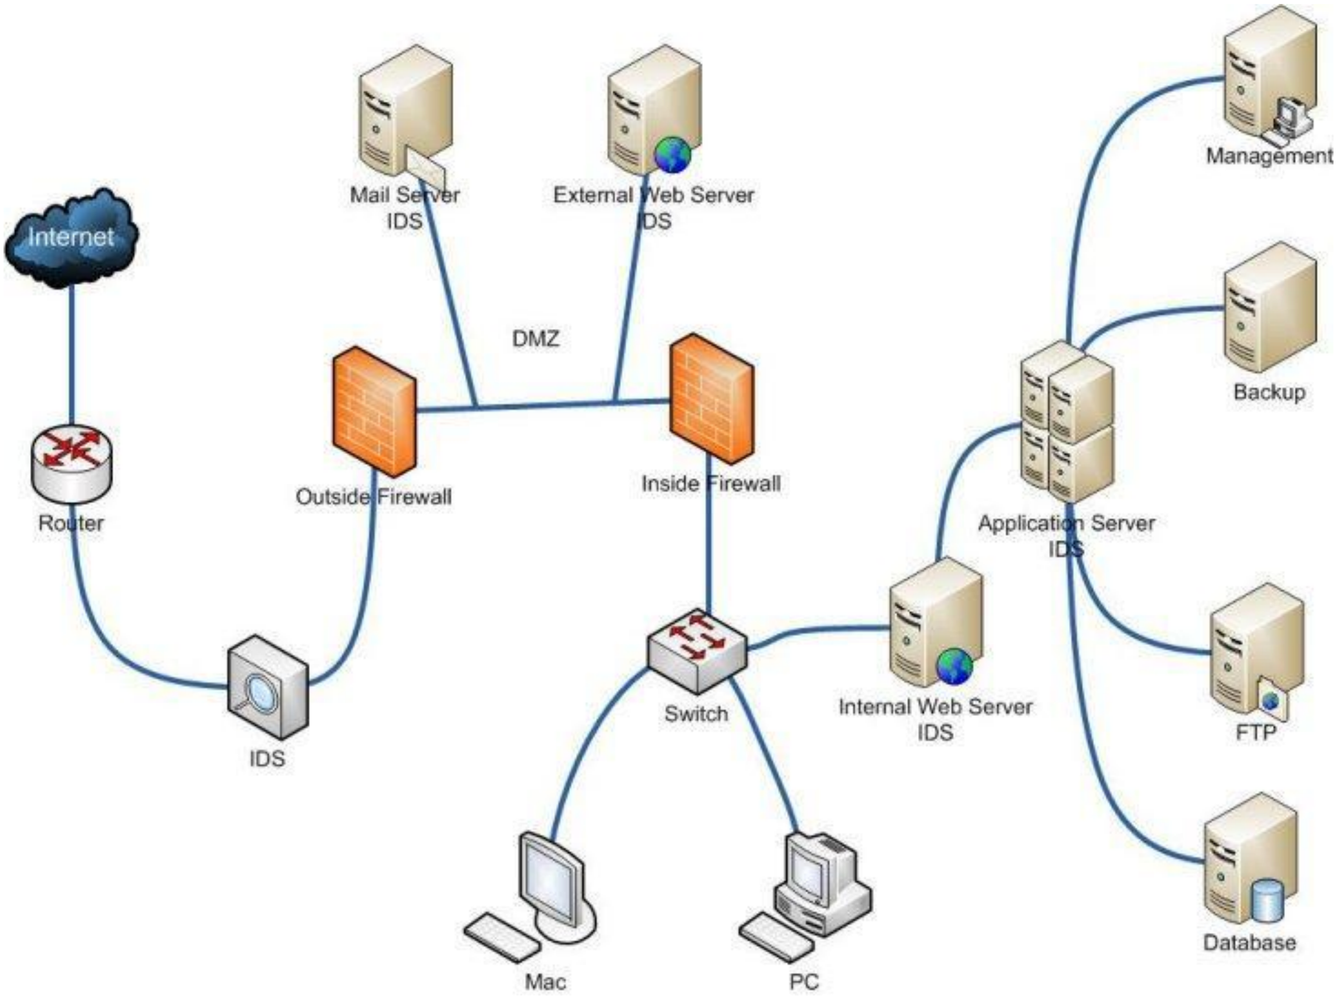

Network Configurations

The focus here has been to demonstrate why security is of such paramount concern to the company. We have outlined the company's electronic assets and the risks which accompany them. We would be remiss if we did not, however, at least offer guidance on how to implement a few of the precautions we have recommended. For that reason, we will briefly focus on protection measures and how they can be configured. Two of the most valuable protections for any network are a Firewall and an Intrusion Detection system.

Iptables Firewall

Most often the first line of defense against malicious traffic is the firewall. Network firewalls can be used to filter network traffic before it can enter the internal network while local firewalls can be configured on servers and hosts to ensure that malicious traffic that has made its way inside the network still cannot harm the host. It is good to have layers of defense to ensure that the network remains protected, even if the first line of defense should fail.

The purpose of a firewall is to define a set of permissible traffic, and deny everything else. This ensures that all regular, day to day traffic can get through, but any other activity will be stopped before it can reach its destination. This also applies to traffic that is attempting to pass through a system on its way to another destinations. Most systems come with a firewall already installed and ready to work out of the box. Even newer versions of Windows include Windows Firewall which offers protection against unwanted traffic and can be configured to be more or less permissible, depending on the needs of the system. Most Linux systems include the iptables firewall by default.

Iptables can be configured to be an extremely effective firewall. What it cannot do on its own, it can be coupled with other programs to accomplish. For example, PSAD (port scan attack detector) can be used to parse iptables logs and look for attacks such as port scans and take action on them if they are discovered. FW-Snort (firewall snort) can also be combined with iptables in order to allow the firewall to block traffic based on snort rule signatures. With so many options available, it can be hard to know where to start. The *Linux Firewalls* book by Michael Rash, describes in great detail not only how to configure iptables, but how to combine it with such tools as PSAD and FW-Snort. A good default iptables policy can be found on the book's website: http://cipherdyne.org/LinuxFirewalls/ch01/. This is a good guide to use as a starting point.

Snort Intrusion Detection System

An Intrusion Detection System listens to all the traffic passing across the network and documents the details of each packet for inspection. Most IDS come with a detailed list of rules which allow them to identify malicious threats from port scans, to viruses, to attempts by hackers to create a back door into the system. Intrusion Detection Systems are valuable tools not only for detecting and recording details about attacks, but for taking steps to prevent them. If coupled with an iptables firewall, Snort can even be turned into an Intrusion Prevention System. It is inevitable that someone will attempt to break into the network. An IDS is the best way to gather the information you need to take action against that and future attacks.

Snort is an IDS which can be used on both Linux and Windows. The configuration is fairly similar on both operating systems. Snort can be configured to identify internal trusted networks and servers. It can be run either in passive mode, sniffing network traffic and returning information about the current network traffic, or it can be run in IDS mode where it will log alerts based on the configuration. Snort can even be configured to email administrators in the event of certain types of detections on the network. The configuration is quite flexible, allowing the danger levels to be defined and the minimum requirements to trigger a reaction for each danger level.

Unfortunately, while Snort and other IDS tools are extremely useful, they are usually quite difficult to initially set up and can be difficult to maintain depending on the source of information they draw on to function. Snort is recommended because the configuration can be carried over quite easily from one operating system to the other. Snort is also quite well documented and the rule set is quite meticulously maintained. A list of guides for setting up Snort can be found on the Snort website: http://www.snort.org/docs/setup-guides/. A detailed tutorial on how to set up Snort and several tools which can be combined with it (including a graphic interface) on Windows can be found here: http://assets.sourcefire.com/snort/setupguides/Snort%20Installation%20on%20Windows%20XP.txt.

Future Considerations

Harding the network and securing the IT environment is no small nor easy task, as we have clearly illustrated. Successfully protecting the company from electronic compromise is no small victory. Unfortunately, the victory will be short lived if the security measures are maintained. It is often said that security systems are not able to function on their own for very long. If left alone, it will not take long for them to become insecure themselves. As new vulnerabilities are discovered and information about them is released to attackers, new threats will appear for the system. With every new software update or release there are new patches to be applied and new risks to consider. Once your security system is in place it is vital that you monitor and maintain it, or the security system may eventually become a risk itself.

The best practice to follow in regards to security is “The Rule of the Three Fold Process”. This is similar to the development life cycle followed for development of new software. The three fold process includes three steps:

1) Implementation
2) Monitoring
3) Maintenance

Once a system has been implemented, it is important to monitor it to make sure it is working properly. You may need to tweak the configuration of the Intrusion Detection System to ensure that it is properly detecting intrusions rather than constantly offering false positives. Intrusion Detection logs, as well as event logs, system logs and authentication logs should be regularly audited and anomalies should be well documented. It is important to learn of an attack as quickly as possible so that proper action can be taken. If an attack is not discovered until months after the fact, all the critical information needed to take action against the attacker may be lost due to the system having been in use for so long after the attack.

Security systems also require maintenance. As new threats are discovered and new patches released, they must be applied to the systems. New definitions which define the criteria that match attacks must be applied to anti-virus software and Intrusion Detection Systems. New plans may need to be made for new security implementations. Changes to the network must be planned for and security considerations taken before any change is made. Changes to the security systems must be even more carefully planned for to ensure that security is being increased rather than put in jeopardy.

Unfortunately, security is a full time job. However, as long as the three fold process is applied and considered for all security concerns, it should not prove difficult to maintain a high level of security.

Conclusion

There are many things that must be considered to ensure a secure IT environment. The services which are being offered and the way in which they will be configured and maintained are all important considerations. The computers which the services will depend on, the computers which employees will use for their work and the networks which connect everything together must be secured to ensure that no loss is sustained by the company. The protection of sensitive information is of the highest concern; the company could be liable for lost or stolen information. It is important to minimize the impact of down time to ensure that the company does not suffer loss of money or productivity. Understanding the threats to the company and the consequences of exploitation of those threats will help to develop an efficient, cost effective security system which will prevent the loss of money or business for the company. While implementing and maintaining the security measures found here may be time consuming and, at times, difficult, it is well worth the effort to prevent major loss or damage to the company or its reputation.

Credits

Threat Risk Analysis
Windows Server 2003: Jaanus Anja

Linux Server: Boris Plotkin

Windows XP Client: Megan Cutler

Network Protection Measures: Daniel Sowinski

Presentation
Windows XP Client: Megan Cutler

Windows Server 2003: Jaanus Anja

Linux Server: Boris Plotkin

Network Protection Measures: Daniel Sowinski

PowerPoint Presentation: Daniel Sowinski

Network Diagrams: Daniel Sowinski

Documentation
Analysis, conclusions and recommendations: Megan Cutler, Jaanus Anja, Daniel Sowinski and Boris Plotkin

Network Diagrams: Daniel Sowinski

Final document layout: Megan Cutler