E-commerce Production Firewalls

A Proper Security Design

1. Executive Summary

The purpose of this paper is to detail the design of a production firewall for an e-commerce company. Companies with websites and other public facing services do not take into account correct security practices for their network. It is important to understand the security needs of protecting their web site and other Internet facing computer systems.

A firewall is the focal point in network and system security. This paper will look at proper firewall standards and best practices, modeled after Cisco SAFE and CERT, for using a firewall in an e-commerce network. Proper DMZ design and the physical placement of the firewall will be discussed. Also, firewall security policy rules, and how best to configure them. Besides normal firewall design, this paper will list other ways to secure the firewall itself, with proper logging and daily backups of the configuration, security audits, and disabling unneeded settings.

This paper will give network administrators a proper guide to securing a network and the firewall.

Table of Contents

Figures

Tables

2. Why use a Firewall

Out of the variety of network devices built to create flexible and stable networks, the firewall is the single most security focused device for network and system protection. The main purpose of a firewall's security feature is to protect and control the access to computer systems.

However, firewalls are capable of doing more than just simple access control:

- "*They can provide a single 'choke point' where security and audit can be imposed*" (Robertson, Curtin, & Ranum, 2004, Section 2.3).
- "*Packet filtering and application proxies. These functions can be used separately or jointly*" (CERT, 1999, Introduction Section).
- "*Firewall—Allows granular control for traffic flows between the management hosts and the managed devices*" (SAFE, Enterprise Campus).

The amount of viruses and hackers running loose on the Internet increases the threat to any company that wish to participate on the Internet; there is no question that security is in any company's best interest. As Parker (2005) states, "*With the ever increasing size of the Internet there is seemingly a never-ending market for the services … hackers can provide*". Security is not only important to control access, but in many cases it's a requirement by board members, stock holders, and Sarbanes-Oxley.

This paper will discuss proper placement of a firewall in the network, creating security policies, best security practices, and conducting security audits. These security recommendations follow such standards as Cisco's SAFE model and CERT's Security Practices.

3. Network Placement

There are many ways to install a firewall into a network. Items to consider are: what networks should it connect to, what type of security to place between these networks, and what services are to be offered to the Internet.

Figure 1 - Collapsed Network shows a company network that is collapsed into one single network. Servers, office, and employee PC's are all located on the same network with no firewall in place between them and the Internet. If any one computer gets infected with a virus or compromised by a hacker, then all computers are at risk and accessible to the virus or hacker with no security restrictions between them. This is obviously unacceptable to any company as work productivity would become non-existent.

Figure 1 - Collapsed Network

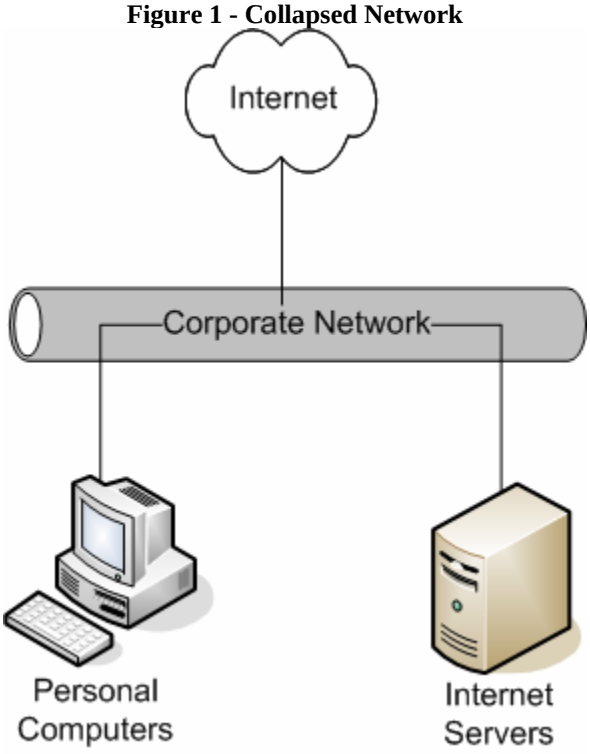

PJ Balsley © 2006

The solution is to segmenting the different computer services with a firewall. As in an E-commerce network, the purpose of the firewall is to connect services to the Internet such as web and email as well as connect the corporate network to the Internet. This segmenting design is referred to as a DMZ, or Demilitarize Zone. Refer to Figure 2 - Segmented Network. By separating the different functions or network services by a firewall each are protected from each other.

Figure 2 - Segmented Network

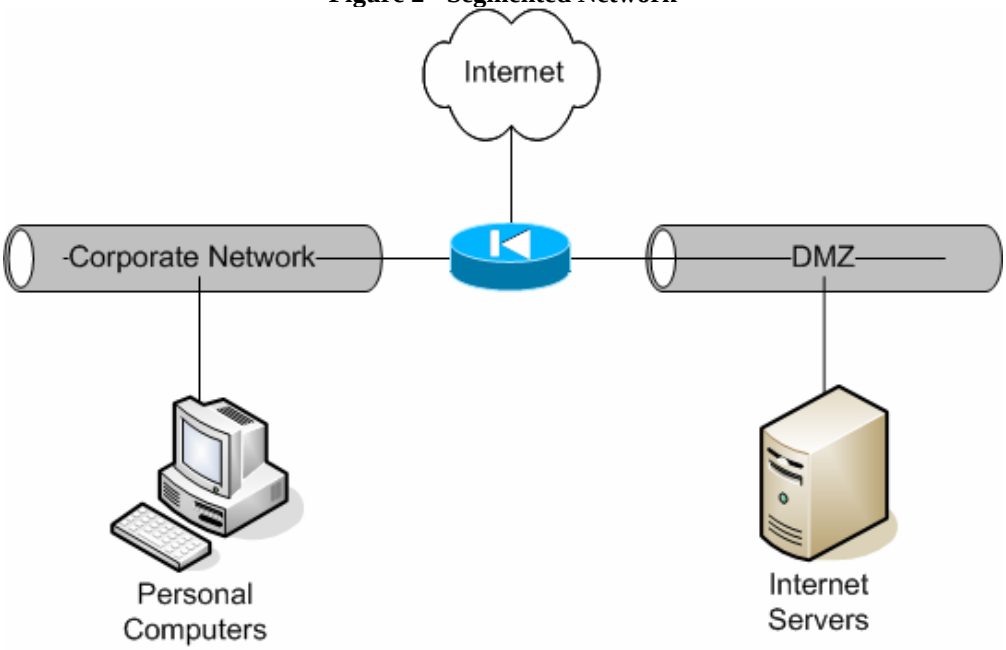

PJ Balsley © 2006

CERT (1999) explains why DMZ's are important. "*In a DMZ network, the untrusted host is brought 'inside' the firewall, but placed on a network by itself (the firewall host then interconnects three networks). This increases the security, reliability, and availability of the untrusted host.*" (Section 'Select the Firewall Topology').

Firewalls typically have three interfaces: Inside, Outside, and DMZ. These three interface functions typically allow the firewall to connect a corporate network to the Internet and place servers in a protected DMZ zone for Internet only access. The DMZ is then denied access into the Inside interface network.

3.1. Inside Interface

The inside or fully trusted interface should be plugged into the main corporate network. This will allow corporate access to the Internet and to the DMZ interface. This interface is considered fully trusted because it will have the least amount of inbound access from the Internet, therefore the least amount of security risk.

The corporate network is where all employee computers, printers, domain controllers, and other corporate based systems would exist. These computer systems normally need access to each other and external Internet connections, such as web browsing, but very rarely would require any incoming connections originating from the Internet.

3.2. Outside Interface

The outside or external interface is connected to the public network, which is the Internet side of the network. This is the un-trusted side, as a security model the Internet is an un-trusted source. The outside interface is where public hosts, from the Internet, will connect to the corporate DMZ services, such as a web site or an email server.

Any access allowed in from the outside interface should be strictly controlled. Robertson, Curtin, and Ranum (2004) suggest that, "*you should consider blocking everything by default, and only specifically allowing what services you need on a case-by-case basis.*" (Section 3.11). This means if the only servers located in the DMZ network are web servers then only allow TCP port 80 in from the Internet.

3.3. DMZ Interfaces

The DMZ interface is designed to separate all Internet accessible systems from the internal corporate network, or inside interface. The DMZ computers are controlled by the company and need to be protected, but since they connect to the Internet and have the potential to be exploited, they are not completely trusted and therefore will not have full access into the internal corporate network. If one DMZ system is compromised then it would not easily spread to the internal corporate network. Refer to Figure 3 - DMZ Example.

Figure 3 - DMZ Example

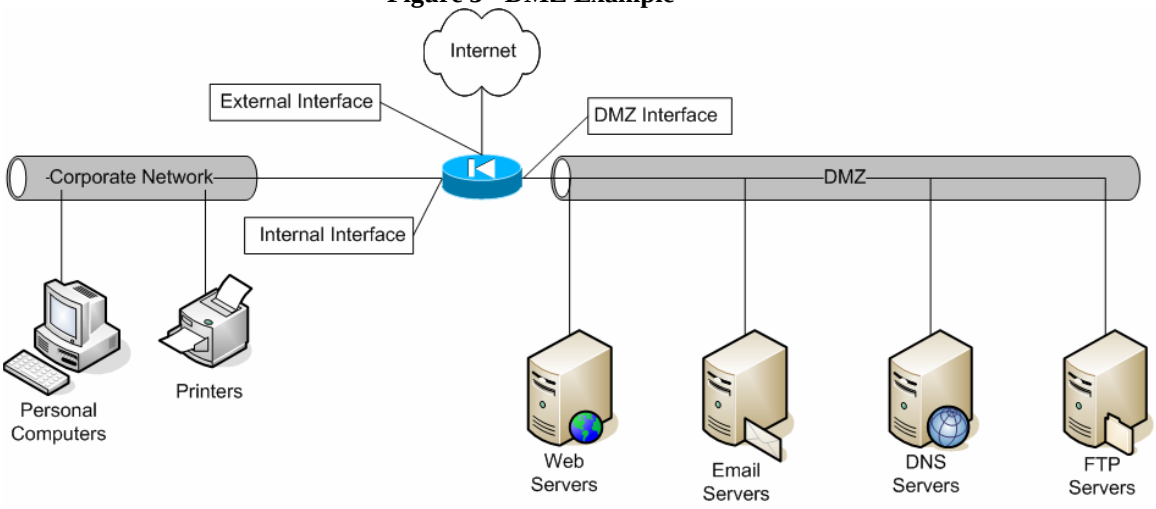

However, a problem with having a single DMZ network is if one of the servers gets compromised, then the hacker has full access to all surrounding servers. This allows all DMZ servers to be at risk from each other. Thus it is becoming common practice to create multiple DMZ networks to segregate the servers from each other giving each of them more protection. Robertson, Curtin, and Ranum (2004) state that, "*If you are running a number of services that have*

different levels of security, you might want to consider breaking your DMZ into several 'security zones'." (Section 3.9).

For example, the network design in Figure 3 - DMZ Example has the need for the following DMZ zones: web, services, and FTP. This will group all web servers into one DMZ, all FTP servers into another DMZ, and the email and DNS into a third DMZ. All three of these services do not need to directly interact with each other and can function independently in their own DMZ zones. See Figure 4 - Segmented DMZ. Now that our Internet services are placed in their own DMZ zones, they will not affect each other.

Figure 4 - Segmented DMZ

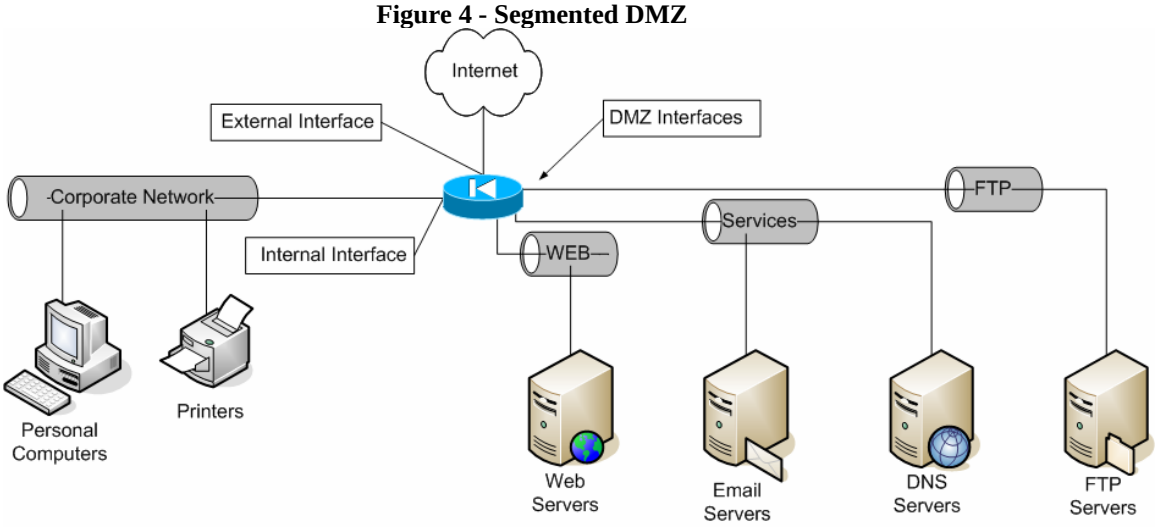

PJ Balsley © 2006

First let's talk about two ways to accomplish this number of DMZ zones. One way is to purchase a firewall device that physically has the number of ethernet ports to assign one per zone. In our example this would require five (5) interfaces. That is three (3) for the DMZ zones and don't forget the internal and external connections as well. Using physical ports is the most secure way to separate the zones, but the draw back is cost and limitation.

Finding a firewall that has five or more interfaces can get expensive and limits your choices in firewall vendors. What about future growth? What happens when more DMZ zones are required? Larger more complex E-commerce networks can have 12 or more DMZ zone requirements. There are only so many interface ports that can be installed in any firewall.

A popular alternative to this is virtual VLAN trunking from one single interface. This allows for greater zone creation and is limited only to the firewall vendor software. Vendors that can do this have limitations starting around 100 interfaces and up. This allows for the virtual creation of network interfaces using IEEE 802.1q tagging. Thus, allows for a low cost and almost limitless solution.

4. Rules and Policies

Firewalls are used in networks as a security barrier or buffer. They are designed to block or pass specific traffic between interfaces. When designing the firewall rules, or policies one must consider what services will be offered to the Internet, to the corporate network, and what other functions will the firewall be required to fulfill.

Rules must be built on a per basis design. Each network will need different ports and services allowed depending on the needs of the business.

4.1. Outgoing Traffic

Outbound traffic is all corporate network traffic going out towards the Internet. Many companies allow full access to the Internet. This puts the least amount of impact on users when using programs to surf the web, check email, transfer files, stream music, and so forth. They believe that people work hard at their jobs and do not wish to put unneeded restrictions on them. Reward the employee with the ability to do anything on the Internet.

On the other side, companies will block all ports only allowing what services are needed such as web and email. They do not let anything through that is not authorized and will even block what web sites employees can view. This is a very strict locked down, but makes for a secure network. Why? Because viruses, spam worms, and exploits can not be sent out to the Internet and embarrass the company or employees. This also prevents employees from wasting time and goofing off by sending personal emails and web surfing for fun.

To fulfill both security and personal activity concerns common ground can be found. To allow employees proper access to the Internet most ports should be allowed out of the network, but some specific services should be blocked. SMTP should only be allowed out from defined email servers, not personal computers. This way email virus originating from email worms on personal computers will not propagate outside the company. Instant messaging ports should be blocked so users can not chat the day away with friends. But instead, setup an internal instant massaging server so employees can conduct business with each other. Again, your corporate policy will dictate how strong or lax the corporate policy should be.

Figure 5 - Outbound Traffic

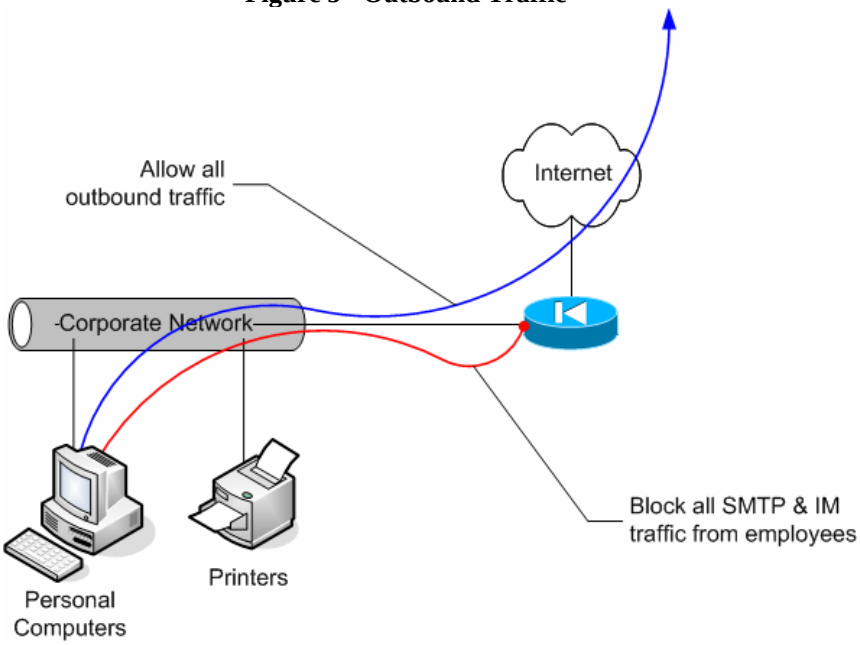

PJ Balsley © 2006

4.2. Incoming Traffic

All inbound connections from the Internet should be denied to the corporate network, unless a very specific application is needed, but if this is the case, the server should be moved to a DMZ network instead of the corporate internal network. This will isolate and keep safe corporate employees from most hackers and worm viruses.

For inbound connections from the Internet to the DMZ servers, only allow the ports needed to run the functions the server is needed for. Email servers should have TCP port 25 allowed to them. Thus a hacker only has one possible point of entry to try to compromise the server with. If all ports where allowed open, then that hacker could try to login via telnet or ftp or other services on the server that do not need to be allowed open to the Internet.

Figure 6 - Incoming Traffic

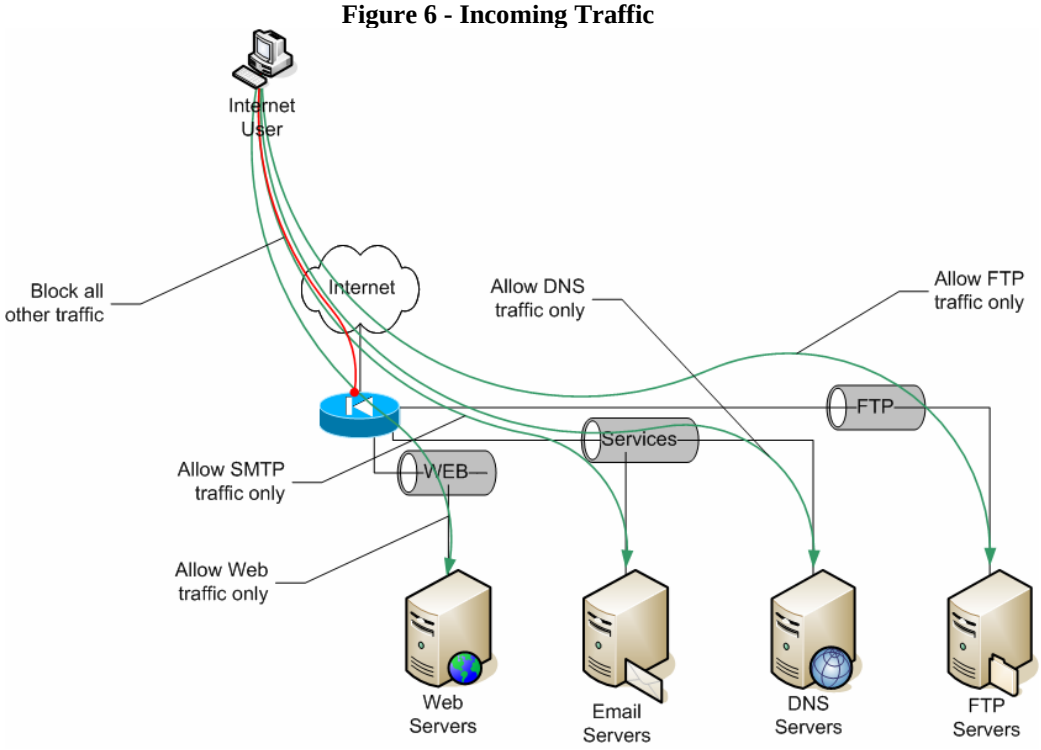

PJ Balsley © 2006

4.3. Corporate office to DMZ

Since the internal network is completely trusted, access to the DMZ network is normally allowed with no restrictions. Allowing full access will allow administrators and developers unrestricted access to do their jobs building, testing, and administrating the functions on the servers. Only specific business requirements would dictate if any ports need to be blocked for special requirements.

Figure 7 - Internal Traffic

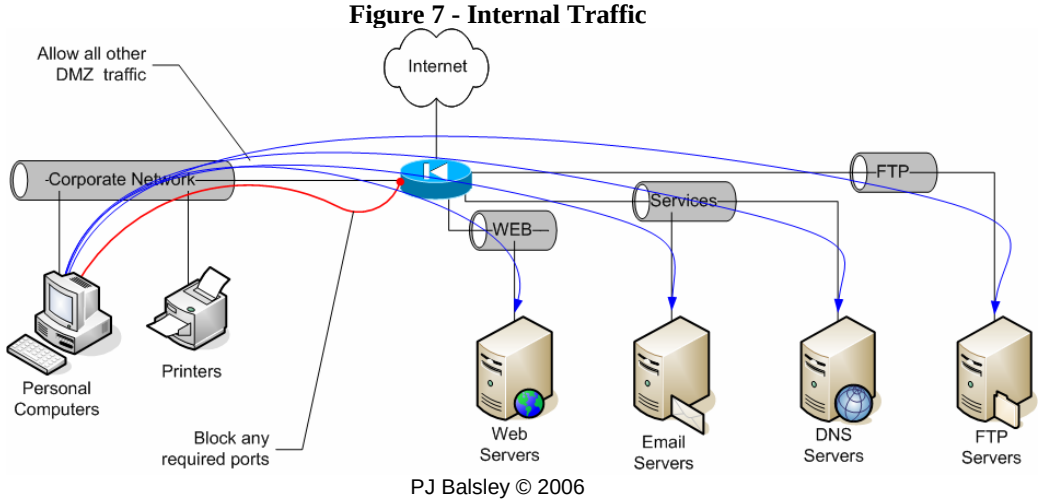

PJ Balsley © 2006

4.4. DMZ to Corporate office

Servers in the DMZ zone should be considered un-trusted. This is because they have inbound connections from the Internet and are targets for any type of virus, worm, or hacker. They may be behind the firewall for protection, but they still are susceptible to attacks. This is the reason that DMZ servers should have limited communication into the corporate network.

Best practice is to deny all connections from the DMZ to the internal network, unless specific applications are needed. Such as, the web and email servers may need access to the DNS or LDAP authentication server located in the internal corporate network. In this case, only open the needed ports between the web and email server IP's and the DNS and LDAP server IP's. This will prevent a compromised DMZ server from getting full access to all corporate computers.

Figure 8 - DMZ to Corp Traffic

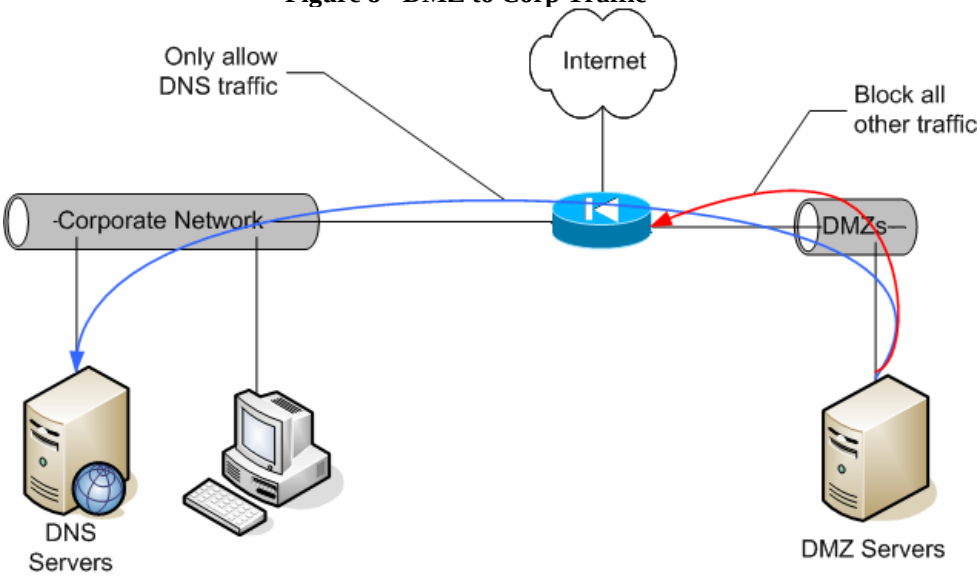

PJ Balsley © 2006

5. Security Practices

Configuring the firewall to protect the network is just the first step in creating a secure network. Many people make the mistake of thinking that because a firewall is in place, the network is 100% secure. The policy rules might be well written, but the security of the firewall itself must also be taken into consideration. There are many simple steps that can be made to help secure the firewall devise such as access restrictions, auditing and logging.

5.1. Access Restrictions

It is important to control the level of access to the firewall. This means strong login information and passwords as well as restricting which internal and external networks are allowed to logon and manage the firewall.

Because the firewall is publicly accessible to the Internet, this allows anyone the potential ability to login and change firewall settings to gain further access into the network it protects. The Computer Emergency Response Team (1999) has stated:

> *If you need to administer your firewall systems remotely, you must use strong authentication and data encryption technologies to prevent adversaries from compromising your firewall systems. The firewall administrator should be authenticated using technologies such as one time passwords or recognized cryptographic protocols rather than using clear text passwords or replayable authenticators. All administrator communications to and from the firewall systems must be strongly encrypted. Consider strongly encrypting any sensitive information (such as passwords, configuration data) stored on the firewall system or on all administrative systems (such as the network management system).*

Not removing the firewall vendors default password or choosing an easy to guess password is one of the largest mistakes to make for password security. In selecting a password, there are some simple rules to follow. Fraser (1997) states them as:

> *They should not be single words in any language, any common, industry, or cultural acronyms, etc. Ideally, they will be longer rather than shorter and consist of pass phrases that combine upper and lower case character, digits, and other characters*. (Section 4.1.3).
>
> Fraser (1997) also list out some password tips (Section 4.1.4):
>
> - *Change default passwords.*
> - *Routinely change login passwords.*
> - *Disable login account if too many failed attempted logins occur.*
> - *Change all passwords if a system is compromised.*

Following these guidelines will help insure that the password is hard to guess, crack, or brute force. An example of a good password would be: *g#tOut40w*. Also, never write down the password where someone could read it. If there are too many passwords to remember, then use a password manager that will save all passwords in an encrypted file format.

Most firewalls, as well as other network devices, allow for restricting login from specific networks. This should be employed to only allow a telnet or SSH login from the networks that have been authorized. These networks could be the netblocks that all network administrators PC's are plugged into, or from specific computers where authorized access is allowed. It is never a good idea to allow external, Internet, network access to manage the firewall. If this is necessary, then only use a static IP address of the computer that will

need access. Make sure this IP never changes like many DSL home providers will.

If the firewall supports SSH as an encrypted login protocol, then always use it over telnet. Telnet is a clear text communication protocol that passes the login username and password over the network and can be easily captured by a sniffer program. Using a secure shell (SSH) will encrypt the login information before connecting to the network device and thus hiding all login information from prying eyes. SSH will also rotate its encryption keys normally every hour, to prevent eavesdroppers from cracking the SSH session and reading your information.

5.2. Backups & Auditing

Another step to maintaining firewall security is daily backups of the firewall configuration. This is for historical configuration records as well as auditing changes in the settings or policy rules. This could be done manually or from an automated program.

One such program that does this well is called RANCID. "*RANCID monitors a router's (or more generally a device's) configuration, including software and hardware (cards, serial numbers, etc)*" (Shrubbery Networks, 2006). RANCID will nightly download and save network configurations, then check the difference between the last backup and the new one. Any changes are emailed to the network administrators for review.

Two main reasons for using a program like RANCID are if the firewall hardware fails and a replacement is needed, the new firewall can be rebuilt quickly with the last days backup configuration. This will cut down on hours of manual policy entering. This also allows for junior network administrators the ability to check firewall policies without giving them direct access to the firewall; they can just review the backup configurations instead.

Second reason, by monitoring the configuration changes, this allows administrators to audit their peers work, as well as keep up with network changes they are not involved with. Most importantly, by checking the difference in configurations for unauthorized access and alterations, you will be keeping an eye out for hackers or other unauthorized access to the firewall.

5.3. Log Data

Keeping all log data from the firewall is necessary to debug problems, as well as track user access and identify intruders. Most firewalls will have a limited buffer of memory to list log data. It also may be necessary to track log information back in time to weeks, months, or years. It will be important to send all log information so a centralized server that is responsible to save all log data and to a backup file system such as tape or CDROM. Fraser (1997)

stresses the importance to save log data. "*Audit data should be some of the most carefully secured data at the site and in the backups. … Audit data may also become key to the investigation, apprehension, and prosecution of the perpetrator of an incident.*" (Section 4.6.4).

The Cisco SAFE model also recommends saving log data, "*The event-logging server is used to consolidate syslog, IDS, and IPS events to a single source. This provides the forensics necessary to identify, isolate, and recover from a host or network attack.*" (Section Design Guidelines). Saving log data on to another host also prevents the possibility that the logs will be changed or deleted to cover tracks.

5.4. Optional Settings

Optional settings on a firewall should not be overlooked. Many of these optional configurations left at the default value could cause security holes or allow for service disruption. Examples of these options are NTP, SNMP, and HTTP Server.

Network Time Protocol (NTP) should always be set to an internal NTP server to sync the date and time on the firewall. This will make sure that logs are reported in the correct time. This is very helpful when comparing logs from different devices to track or audit problems.

Simple Network Management Protocol (SNMP) community strings should never be left at the default "public" name. SNMP can be used to find out a variety of information about the firewall and its function, as well as make changes to the policy rules and configuration. SNMP can be a very powerful tool and needs to be secured or disabled if not used. Cisco SAFE (n.d.) stresses, "*SNMP should be treated with the utmost care*". (Section Secure Management and Reporting).

Many firewall vendors use a HTTP web server as an optional configuration utility. Historically web servers can be exploited on these devices to gain unauthorized access, delete files, and crash the hardware. Unless this is the only form of configuration utility, it is recommended to turn off the web server on the firewall.

Lastly, turn off any unneeded services or running applications. If a service is not required for the firewall operation or needed to manage the device, turn it off or disable its access to the network. The Cisco SAFE (n.d.) model lists some simple steps to secure a device:

- *Locking down Telnet access to a router or using more secure methods such as SSH*
- *Locking down SNMP access to a router*
- *Controlling access to a router through the use of TACACS+*

- *Turning off unneeded services*
- *Logging at appropriate levels*

6. Security Audits

Security audits are an import part of monitoring the firewall's effectiveness and relevance to the current network security needs. Auditing on a regular basis will reduce potential security holes and threats. At least two types of audits should be conducted: network scans, and configuration review.

6.1. Scans

Network scans reveal what ports or services are available over the network and analyzing these services for reported vulnerabilities. This is one of the first things that a hacker would do to find out how to break into a network. Conducting the same scans and fixing any reported problems should be a high priority. One such network analyzer is a program called "Nessus", www.nessus.org.

6.2. Configuration Review

A quarterly review of the firewall policy for changes and rules should be conducted. This is the time to clean up the configuration for services that are no longer needed, implement new policies, and review the last 3 months of changes. This review should be done with all network and security administrators to assure that no configuration is misunderstood or interpreted.

7. What won't a Firewall do?

This question may be the most important one for securing a firewall. Knowing the firewall's weaknesses give the administrator the knowledge and ability to address, identify, and fix or create work-a-rounds to the firewall weaknesses.

Robertson, Curtin, and Ranum (2004) list items that a firewall can not protect against: (Section 2.4)

- *Firewalls can't protect from attacks that do not go through the firewall.*
- *Firewalls can't protect against access already allowed.*
- *Firewalls can't protect from internal attempts.*

As listed, one of the problems for firewalls is if a service is allowed through, lets say HTTP for a web server, and the web server is attacked via a web exploit the firewall will allow such an attack, as it looks like normal web traffic. Because of these types of problems many firewall vendors are beginning to build antivirus and intrusion protection proxies into the firewalls. This will allow the firewall to identify the difference between real valid web traffic and a malicious attack.

Unfortunately, many of these features are still new and not always effective in firewalls today. Therefore, there are plenty of antivirus and intrusion protection devices that can be used in the network along side the firewall to protect against these possible firewall weaknesses.

8. Firewall Compliance Checklist

The Compliance Checklist is a list of control measures that test the firewall/network design for proper compliance to this document.

Table 1 - Compliance Checklist

Compliance Control Measure	Pass	Fail
Created needed DMZ interfaces		
Block all inbound corporate connections from the Internet		
Set outbound corporate traffic as needed		
Setup required inbound connections for DMZ servers		
Changed vendor default passwords		
Changed or set SNMP community strings		
Turned on NTP		
Restrict access to specific IP/netblocks		
Enable SSH / disable telnet		
Backup running configurations nightly		
Check backup configurations for changes nightly		
Setup syslog logging		
Disable HTTP server		
Disable un-needed services		
Setup external authentication i.e.: Tacacs+		
Quarterly configuration audit and review		
Network security scanning schedule		

9. Conclusion

Security is a concern for every company and more so for companies that wish to place publicly accessible computers on the Internet. Many network and security administrators are unaware of proper security configuration practices and this paper will help them abide with Cisco SAFE and CERT recommendations. Following these recommendations will create a more secure network design to protect the company's important computer systems.

10. References

CERT. (1999). Design the firewall. Retrieved Jan 25, 2006, from http://www.cert.org/security-improvement/practices/p053.html

Cisco. (n.d.). SAFE Blueprint. Retrieved Jan 19, 2006, from http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns128/networking_solutions_package.html

Fraser, B. (1997, September). Site Security Handbook: Request for Comment 2196. Retrieved Jan 19, 2006, from http://www.rfc-editor.org/rfc/rfc2196.txt

Parker, D. (Jan 18, 2005). The Convergence of Hacking and Security Tools. Retrieved Feb 25, 2006, from http://www.windowsecurity.com/articles/Hacking-Security-Tools.html

Robertson, P.D., Curtin, M., & Ranum, M. (2004, July 26). Internet Firewalls: Frequently Asked Questions. Retrieved Jan 19, 2006, from http://www.compuwar.net/pubs/fwfaq

Shrubbery Networks, (n.d.). RANCID - Really Awesome New Cisco conflg Differ. Retrieved Feb 25, 2006, from http://www.shrubbery.net/rancid/