Global Information Assurance Certification Paper

FTP Server Security Strategy for the DMZ

Joseph M. Adams
SANS Security Essentials GSEC Practical Assignment
Version 1.2d

Introduction

Recently, the company division has made the decision to utilize the Internet as much as possible for data/file transfers to and from our external customers. This is a change in strategy from previous projects, which relied on private networks and private data communication lines. With the prevalence of business presence on the Internet, this strategy is a good way to reduce telecommunications cost for companies and their vendors and business partners. Other reasons include the leveraging of the infrastructure across multiple customers and reducing the number of staff technical support hours. However, before any major change in strategy and policy, customers and management want to be ensured that proper security practice and procedures are utilized to protect the customer and internal systems and data.

It was determined, that in order to meet this goal, the technology group needed to redesign the existing infrastructure and create an architecture in which to build upon. This paper will present how this change of corporate direction demanded a need for tighter security, what design was chosen and how it was implemented.

Security Program

The first area that needs to be addressed is the Security Program. Prior to any changes in strategic direction, new technology or ideology, the security program should be updated or created. The reason for this is to identify and address any issues associated with a new technology BEFORE a real issue occurs that could threaten the security of the infrastructure.

Any major shift in technology such as Internet file transfers or web server implementations should not be taken lightly. These strategies need to be part of an overall security program. A good security program should provide the appropriate balance between prudent levels of control and ease of access. To maintain the confidentiality, integrity, and availability of information resources, a security program should implement a well-balanced, risk-based information security program. Only through understanding and adhering to this type of program and its associated policies, directives and standards will risk be minimized and resources protected from external/internal threats. The security program needs to address areas such as audit, administration, platform, application and network security, incident response and overall user policies. Prior to the implementation of the Internet FTP servers, the current security program needed to be updated to address the new strategies. For example, the

Unix platform security policy needed to be updated to account for a server in the demilitarized zone (DMZ).

By implementing the changes to the existing security program, it is demonstrated to customers that their interests, as well as the company, are served.

The DMZ for the FTP Servers

The first part of the design is the DMZ. An environment needs to be created to house the FTP servers and provide protection for the company specific infrastructure.

A DMZ is a network between the protected network and the external network. (See Figure 1) The DMZ provides an extra layer of security and is a good place to locate hosts that provide web, ftp and smtp services. Servers that house these types of protocols are referred to as bastion hosts. Mr. Craig Hunt, author of "TCP/IP Network Administration" defines a bastion host as follows.

> "A bastion host is a secure server. It provides an interconnection point between the enterprise network and the outside world for the restricted services. Some of the services that are restricted by the interior gateway may be essential for a useful network. Those essential services are provided through the bastion host in a secure manner. The bastion host provides some services directly, such as DNS, SMTP mail services and anonymous FTP."[1]

Firewalls are positioned on each side of the DMZ. The outer firewall is more of a "screening" firewall; in that it will block certain protocols, but let others through that are allowed in the DMZ. The outer firewall is allowing FTP, port 25 traffic (as well as other protocols such as http, port 80) into the DMZ. The inner firewall, which is protecting the internal network, is denying these protocols from entering the internal network.

Machines located in the DMZ have the greatest exposure to potential attacks. The security implications of any application-level or protocol-level connectivity between external machines and internal machines must be carefully considered, since improper implementation may result in increased vulnerability of internal networks to external attack.

The advantage to placing servers in this zone is to prevent any unauthorized traffic from entering the Internal Network. Once the DMZ is designed and built, the Intrusion detection system devices and the FTP servers can be positioned. The benefits of a DMZ design are noted by Mr. Chuck Semeria of 3Com (Mr Semeria used packet-filter routers instead of firewalls). The following quote is taken in full from his whitepaper titled "Internet Firewalls and Security A Technology Overview"[2]

> "There are several key benefits to the deployment of a screened subnet firewall system:

- An intruder must crack three separate devices (without detection) to infiltrate the private network: the outside router, the bastion host, and the inside router.
- Since the outside router advertises the DMZ network only to the Internet, systems on the Internet do not have routes to the protected private network. This allows the network manager to ensure that the private network is "invisible," and that only selected systems on the DMZ are known to the Internet via routing table and DNS information exchanges.
- Since the inside router advertises the DMZ network only to the private network, systems on the private network do not have direct routes to the Internet. This guarantees that inside users must access the Internet via the proxy services residing on the bastion host.
- Packet-filtering routers direct traffic to specific systems on the DMZ network, eliminating the need for the bastion host to be dual-homed.
- The inside router supports greater packet throughput than a dual-homed bastion host when it functions as the final firewall system between the private network and the Internet.
- Since the DMZ network is a different network than the private network, a Network Address Translator (NAT) can be installed on the bastion host to eliminate the need to renumber or resubnet the private network."[2]

Figure 1 depicts the computing environment that was designed for the file transfers. Note: This DMZ is not the environment that houses mail or web services, therefore, these services are not identified in the following diagram.

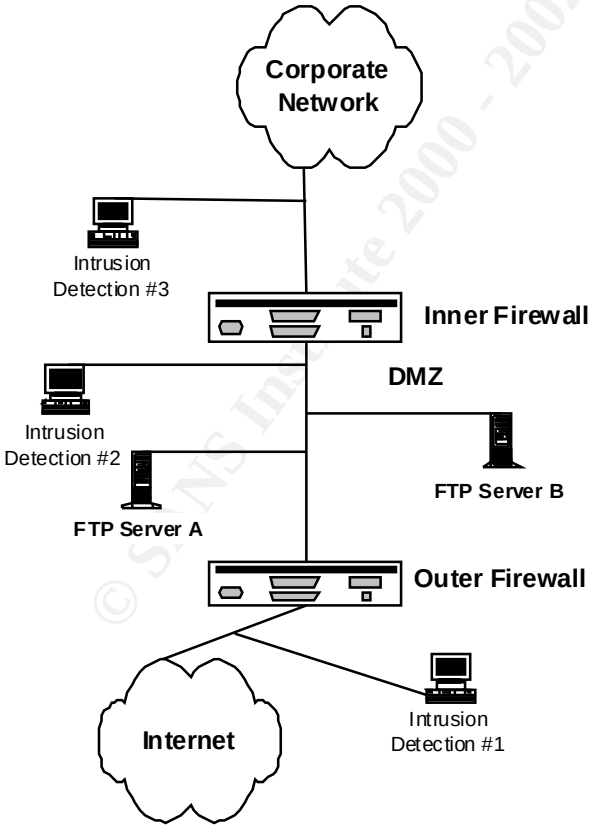

Figure 1

Securing (hardening) the FTP Servers

The UNIX server configuration, used to host the FTP service, is defined in a policy under the Security Program. The policy dictates that the Operating System will be installed and patched in compliance with vendor requirements and system security policies. The Operating System will be the latest supported version offered by the vendor, provided that it is stable. The Operating System should be kept current with patch releases because installing the latest patches and releases of key system executables is essential to maintaining the security of the system. The hacker community is actively searching for new vulnerabilities to exploit, and a vendor announcement of a new vulnerability will often find many that will attempt to exploit it soon after it is announced.

Before the FTP servers can be put into service, each host machine will be installed and configured utilizing the concept of least privilege and service and following the above mentioned policy. The UNIX administrators are responsible for the hardware and software installation/configuration of these machines and they follow stringent guidelines to harden these machines. The policies and procedures that they follow were developed under the Security program.

The Operating System (the FTP servers are running UNIX, specifically Solaris 8.) is installed from the vendor-supplied media. The operating system that is pre-installed on the machine will be removed from the system during the installation process. This way, unneeded operating system packages will not be installed, also, the administrator is sure of the state of the operating system when completed. The operating system is installed using the default core installation option. After the operating system is installed, all non-required packages are removed. Also, any additional required packages, that are not part of the core installation are added at this time.

The hardware should not be connected to the network (private or public) until the system is secured. Each FTP server will not offer or accept more services than absolutely necessary. Name services, such as Network Information Service (NIS) will not be used on the server. TCP/IP services not needed will be disabled by commenting out the entry in the in the inetd.conf file. For example, these services will be removed: name, exec, comstat, talk, finger, uucp and tftp. No unnecessary applications, such as compilers and databases will be loaded onto the servers.

After the operating system is installed, the configuration process begins. This process consists of reviewing system file permissions to ensure they have the proper owner, group and world permissions. The .rhosts and hosts.equiv file permissions should be owned by root and the file permissions should be set to 600. 600 is interpreted as owner (root) has read and write, while group and world have no rights on these files. The next step is to turn on system auditing.

Each user account will be established according to defined procedures and the user security policy standards. Further, only necessary user accounts will be added into the system. The user security policy defines items such as minimum password length and how often password should be changed.

Lastly, each machine will be subject to daily audits. Audit logs and management reports are essential to identify user activity and detect violations. The systems are placed under audit to identify patterns of suspicious or unusual activity. For example, activity attempts, password guessing with or without reaching the threshold, logon attempts, attempts to access critical files and to monitor critical file permissions. Other items reported by the audit are system file dates, which could indicate edits to a file, changes to security parameters, security profiles, including password resets. System modifications performed by administrators or any other users. The audit reports on activities/changes performed to security parameters by privileged commands, transactions or programs. Appropriate file protection is critical to the ongoing confidentiality, integrity and availability of the computing environment

The system administrators in conjunction with the Data Security team will setup, configure and monitor the servers for unauthorized activity. It is the Data Security team's responsibility to review the system audits on a daily basis and report on all unusual activity. The activity report is sent to the administrators and they must identify their actions, to help eliminate normal administrative activities from appearing as questionable actions.

Backing up the FTP Servers

A good backup (and recovery) strategy is crucial to a security program. A strategy was developed to ensure that computing platforms are backed up on a regular basis, that backup data and software are stored securely offsite, and that adequate recovery plans are in place. The strategy is explained below.

The strategy that is used is daily backups of all critical systems. The FTP servers, for the purposes of this paper, are considered critical. Each critical server generates a full back up each evening. The back up media is then sent to an off-site archive and storage facility. Off-site storage, which is part of the overall security program, is utilized in the event the computing facility experiences a disaster. This will ensure that systems may be restored from recent archives.

The media is archived using three classifications, daily, weekly and monthly. The daily backup is archived for a two-week time frame. The weekly backup is archived for a 52-week time frame. The monthly backup is a permanent, which means that this media is archived indefinitely, for the life of the server. Once the media reaches its archive expiration, the media is recycled for further use. For a server restoration, the weekly backup will be recalled from storage and all daily

back ups for that week. The system will be restored from the weekly backup first and then the consecutive days following up to the current day.

The schedule for backups is as follows. Saturday through Thursday, a daily backup is performed. Every Friday, a weekly backup and every first Friday of the month, the Monthly backup is performed.

As soon as the server installation is complete, the backup software is installed. Once this is complete, the server is backed up using the Monthly strategy as described above. This backup can then be used in the event of a total system failure. Once the server is in production, the normal back up schedule will commence.

Why backup an FTP server in the DMZ? The FTP servers will be backup for two reasons. The first is to capture system logs and data, but second and more important for this implementation, is to maintain the backups in the event the system is compromised, so that unauthorized activity can be captured. System Administrators and forensic experts can utilize these backups to reconstruct the system prior to the incident to help in the investigation of the illegal activity and subsequent legal action against the intruder(s).

It should be noted that the backup jobs are running within the protected network on different servers with stand-alone tape drives. The backup data is traversing from the DMZ, through the firewall, to the backup server. The reason for this configuration is cost savings by not having to purchase stand-alone tape drives for the FTP servers. However, this does introduce a vulnerability by having unencrypted data in the DMZ network, which could be viewed by an attacker that is sniffing traffic in the DMZ. Measures were taken to address this issue, however, the resolution will not be shared in this paper due its sensitive nature.

What are Intrusion Detection Systems (IDS)

Intrusion Detection Systems come in two flavors, host and network. Host Based IDS looks at system files and verifies that they have not been modified. Such host-based products include Tripwire and Axent's ITA. Network based, which monitor all packets on a network segment, compares the packets to known attack patterns and raises alerts to those that look suspicious. Such network-based products include RealSecure's Network Sensor.

Network and Host based intrusion detection systems complement each other. Network based is very good at providing early warning signs of attack. Host based intrusion detection systems provide a confirmation to the success or failure of an attack. Also, host based intrusion detection systems provide system specific data such as user name and file name during an attack to help determine the perpetrator. Further, by implementing network and host based intrusion detection systems (along with the firewalls), multiple layers of protection are

provided. The reliance upon a single defense system is avoided and this becomes prevalent in the event a protection layer is compromised.

Network Based Intrusion Detection

The Network based devices are located in each environment; unprotected, DMZ (quasi-protected) and protected. The Network Based IDS is installed on the indicated hosts. Each host has a network adapter card that is running in promiscuous mode. Promiscuous mode will allow the capture of network packets so that the IDS can filter for known attacks.

The devices are physically located in each of the environments in order to monitor each network segment. An unauthorized user will take advantage of any access point in the network, therefore, it is prudent to have a device in each network segment.

Another reason for a network device in each segment is to monitor and ensure that any attacks originating from the previous environment are killed and not allowed through the next level. For example, if an attack originates in the unprotected zone, Intrusion Detection #1 should receive the data, compare the signature to the known attack patterns and take the appropriate action. However, if the Intrusion Detection #1 is unresponsive, the attack could get through to the DMZ. In this case, due to Intrusion Detection #2, the attack should be detected before it can cause any damage to the devices within the DMZ. Further, Detection #2 can also validate the firewall configuration.

Each of the Intrusion Detection devices (#1, #2 and #3) will report back to a central console. The console will be in the protected network. A second console should be in place to serve as a backup to the central console. Also, if the DMZ and private network are in different geographical locations, it may make sense to employ multiple consoles. The console will serve two purposes, configuration and monitoring. The console will allow for the real-time monitoring, data gathering and report generation for the security administrator. In the event an attack is underway, the detection devices will send an alarm to the console.

Communication between the console and detection devices will be encrypted to prevent unauthorized access to the data stream. The console for authentication purposes always initiates the communication channel before issuing requests or retrieving data. This will assist to avoid any attacks, such as spoofing. Spoofing is when an intruder sends messages to a computer indicating that the message has come from a trusted system, in this case the console or Intrusion Detection device.

Host Based Intrusion Detection

Once the FTP servers are setup and secured, the host-based intrusion detection services will be installed and configured. This service runs as a daemon process

on the server. The daemon is configured to check the system logs and compare against certain signatures. If a match is found, the console will be notified immediately and appropriate action will be taken.

Each FTP servers' intrusion detection service will report back to a central console. There will be a backup console in the event the primary is down. The console will be located in the protected network.

File Transfer Process

Now, the FTP servers are open for business. With the infrastructure in place, the file transfer process may begin. Several developers have written scripts to move files to and from the FTP servers. The FTP server administrators have created customer user ids and passwords and the customers have been given their information. They are now able to send or receive their files and data via the FTP server.

Future considerations for data/file security could be file encryption using PGP (Pretty Good Privacy) or a firewall to firewall Virtual Private Network (VPN). Both of these technologies would provide for more security for the data, as the data would be encrypted as it passed through the Internet space. Lastly, Secure Shell may be configured on the FTP server to allow for secure copies (scp) from the FTP server through the inner firewall to the internal server. This is desirable as data passes through this secure zone.

Conclusion

It should be noted that there are probably other methods to architect the above configuration. This configuration was chosen as the best method to accommodate customer requirements and the corporate budget. This being said, a solid security program should still be at the forefront of any such implementation prior to the first production day.

References

[1]Hunt, Craig. TCP/IP Network Administration Second Edition. O'Reilly & Associates. 1992

[2]Semeria, Chuck. Internet Firewalls and Security A Technology Overview http://www.linuxsecurity.org/resource_files/firewalls/nsc/500619.html

Zwicky, Elizabeth D., Cooper, Simon and Chapman, D. Brent. Building Internet Firewalls. Second Edition, O'Reilly Associates, 2000.

Loshin, Peter. Intrusion Detection, ComputerWorld April 16, 2001

Shipley, Greg. Watching the Watchers: Intrusion Detection. Network Computing, November 13, 2000

Internet Security Systems. Network and Host-based Vulnerability Assessment: A guide for information systems and network security professionals http://documents.iss.net/whitepapers/nva.pdf

Orebaugh, Angela. Securing Solaris, October 2, 2000 http://www.sans.org/infosecFAQ/unix/sec_solaris.htm

Fraser, B. Network Working Group Request for Comments: 2196, September 1997 http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2196.html

Velasco, Victor. Introduction to IP Spoofing. November 21, 2000 http://www.sans.org/infosecFAQ/threats/intro_spoofing.htm