



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Technology, Topology and Processes for Corporate Asset Protection During an Internet Attack

GSEC Practical Assignment Version 1.3

Scott Duncan

March 17, 2002

The Internet has become a critical business infrastructure for many corporations. Interruptions in Internet access are considered to be emergencies involving loss of revenue and an added business expense which reduces profitability. Along with this trend of dependence on the Internet, there is an escalating trend of attacks on corporate assets including computer hardware, software, network assets, critical business databases and sensitive internal communications. These attacks come in a variety of forms including viruses and worms, denial of service attacks, and network penetrations. This paper will discuss the implementation of key tools and processes that can limit the damage caused by these attacks and show that it is possible to protect against the unknown.

Example Scenarios

An example from the world of worms and viruses would be the infamous Nimda worm. The first three days of this attack were a very chaotic period for the worldwide Internet community. Many IT staffs, network administrators and security engineers were faced with a storm of attacks against web servers on the perimeter, workstations on the interior and some strange effects on routers and even printers. Some corporations with precautionary measures in place were able to react quickly and minimize the damage. Less prepared companies saw their networks choked, servers locked and workstations unusable within an hour of the initial attack. The damage inflicted by these attacks was reported by E-Commerce Times to be \$590 million [9]. This damage could have been reduced by having processes such as incident response teams in place and security policies to back them up.

Companies who are unfortunate enough to experience a distributed denial of service (DDoS) attack will be familiar with the eerie “silence” that occurs in the initial stages. If a perimeter router has been hit with a DDoS attack such as fraggle[1] or smurf[1], the corporation’s customers and partners suddenly see none of the web sites, e-stores and on-line databases that keep the flow of business moving. Mail servers are not able to send or receive email. The IT staff knows something is wrong but are not able to respond fast enough to avoid a major business interruption, first because they can not identify the problem or secondly they succumb to the chaos in the initial stages. The increasing trend in the sophistication of DDoS attacks along with the automation tools that are available to build attack networks is alarming [2]. Technology such as firewalls, network intrusion detection and network monitoring can be powerful tools for reducing the risk associated with DDoS attacks.

A frightening example scenario in the post-September 11 world would be an electric utility whose network staff has noticed a huge spike in the network traffic going into their network from the Internet. Not knowing how to respond, one system administrator begins to install a packet sniffer to determine the source of the spike. After he retrieves some old

hardware out of the back room, installs some sniffing software, and plugs into the perimeter he finds out that the traffic is coming from a Middle East or an Eastern Europe address. At the same time, an electrical engineer is having some problems using the SCADA system that controls a portion of the electric grid that his company services. The system was previously upgraded with the latest software which provides some wonderful new browser based management features. The electrical engineer calls up IT support who is not aware of the network spike or the suspicious source. Unfortunately, these first two hours of confusion and poor communication do nothing to stop the power being turned off for a whole regional area by a foreign hacker who was able to get browser access to the SCADA system.

Although this scenario has not yet happened in the wild, many electric utilities are starting to realize the to address this possible scenario. The SCADA systems that power utilities use to manage their portion of the power grid were once thought to be too complicated and require too much specialized training to be a viable tool for Internet hackers to attack power infrastructure. With the advent of new web management tools for these SCADA systems and easy web access to training materials and reference manuals, the threat of exploitation of these systems by Internet attackers has become significant [4]. Implementation of key security technologies, good network design and sound security processes are the only way to protect the nation's power infrastructure from this threat.

Technology

Although the task of reducing the risk to corporate assets from Internet attacks is difficult, the good news is that the information security industry has many technology offerings that can be used in the design of a good defense posture. The critical technology components that can be used for early threat detection and response are described below.

Firewalls

Firewalls are important components for a good defense and, when carefully placed in both the perimeter and the internal LAN, can provide *defense in depth*. The concept of defense in depth involves having layers of defense which can fail and not have a catastrophe because the next layer of defense is in place and ready to repel the threat. Placement of firewalls in a good defense posture is further discussed in the Topology section of this paper.

Firewall logs are a valuable monitoring resource when looking for early signs of an Internet attack. For example, if the perimeter log shows that a large number of scans on port 22 (secure shell) are coming from a particular host or network, a rule can be placed in the firewall rulesets that denies all access from the host or network that is performing reconnaissance on the perimeter. By "blacklisting" the sources of the early reconnaissance the Internet attack that may have followed the initial probing can be subverted.

Although firewalls provide a first layer of defense and also provide intrusion detection through logging, it is important to note that reliance on this component without good

security policy and network design is a formula for disaster. In addition, firewall policies can become very complex and should be reviewed on a regular basis [7].

Network Intrusion Detection Systems

Network Intrusion Detection Systems (NIDS) are capable of detecting a large number of Internet attacks and probes by placing sensors on the perimeter of the network [5]. A NIDS places a minimal amount of additional load on the network and can be configured to operate without an IP address in what is known as “stealth mode”. While not impossible to detect when in stealth mode, the attacker that does not immediately find the NIDS may make the mistake of assuming one is not there.

Even though these systems are best at detecting known attacks, they are very valuable when faced with a new Internet attack that uses an array of known threat vectors with the addition of some new tricks. Examples of this were the Code Red and Nimda worms which use a combination of new email viruses and old web server vulnerabilities to spread. If a corporate network had a NIDS system at the perimeter when these worms were first released, alerts would have been triggered for the known attacks and quick analysis of the NIDS alerts and web server logs would have generated enough information to formulate an effective response.

There are a variety of NIDS offerings from both commercial and open source providers. The open source offering of the Snort NIDS performs very well in comparison tests and is available on the Internet at <http://www.snort.org>.

Switches with Port Monitoring Capability

Many older network designs use hubs as a distribution system because they are less expensive and easier to manage. The problem with only using hubs is a hacker who has penetrated the perimeter can quickly find a vulnerable machine on the network and install a network sniffer. Even if the penetration is discovered early in the attack, login credentials may have already been sniffed off of the network and used to crack critical workstations and servers. Using switches makes it harder for a hacker to set up network sniffing which may give a company enough time to stop the hacker before the damage or theft has occurred.

Implementing the Network Intrusion Detection Systems discussed previously on a switched network can be done by placing the sensors for the NIDS on a port that has been configured to monitor the network traffic from a collection of other ports.

Network Monitoring Technology

The process of early detection of Internet threats includes watching for suspicious patterns and anomalies in the behavior of the network and the responses of hosts and network components. Network monitoring technology is invaluable in looking for these patterns in network traffic. Of particular interest are spikes in bandwidth usage that cannot be explained by normal business activity or that occur during non-business hours. Unusually high traffic from the Internet directed at the firewall could indicate a Distributed Denial of Service Attack (DDoS) is beginning. Traffic to and from a web

server that is much higher than normal could mean the web server is under attack by a new worm or has already been infected and is trying to infect others.

There are many network monitoring solutions available including the popular open source project known as Multi Router Traffic Grapher (MRTG) available at <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/>.


Virus Protection

Anti-Virus software is critical to protecting corporate assets but can be difficult to keep current on large networks. By using anti-virus software that has centralized management tools, the IT staff can keep the software and virus signatures current without having to constantly visit each workstation and server. Implementing anti-virus software that does not significantly reduce the performance of the client workstations reduces the motivation of the user to disable the anti-virus software. Centralized reporting capabilities provide an excellent source of information when trying to determine the movement of a worm or virus inside of a network.

© SANS Institute 2000 - 2002, Author retains all rights.

Topology

The topology of a network is a very large factor in how well corporate assets can be protected from Internet threats. Once again, the concept of *defense in depth* should be used in deciding what technology components are placed on a network and where to put these components. The following diagram shows a typical one-layer network design that is used by many small to medium size companies.



The problems with this topology are:

1. There is no firewall at the perimeter.
2. All workstations and servers are on the same subnet.
3. Hubs are used as a distribution system.
4. The server hosting the Internet web site for the company is located on the same subnet as the entire corporate network.

5. The same server is used as a mail server and a web server.

While the edge router can be configured to filter some of the traffic from the Internet, it usually does not have the hardening, logging and configuration options to provide adequate protection or early response when not coupled with a firewall. If there is not enough logging information or the ability to quickly reconfigure the router during the beginning of an Internet attack, the only option may be to disconnect the corporate network from the router. Because of the network design, the whole company would lose Internet access.

There are many problems with having all of a company's workstations, business servers and development servers on the same network segment. An Internet attack such as a malicious worm can infect one workstation and then spread throughout the whole network infecting workstations and critical business servers. This was a common occurrence for sites infected with the Nimda worm. Because of the network design, there was no way to protect critical corporate assets such as business databases and development servers.

Web servers are a very vulnerable part of any network. Notice in this topology that if the web server is under attack it can only be quarantined by disconnecting it from the network. Also there is no mechanism for monitoring the traffic to and from the web server which means there is no way to get an "early warning" that an attack is in progress. If a worm infects the web server, the whole network is at risk of infection. If a hacker compromises the web server he has easy access to the rest of the network and can quickly move to another base of operations.

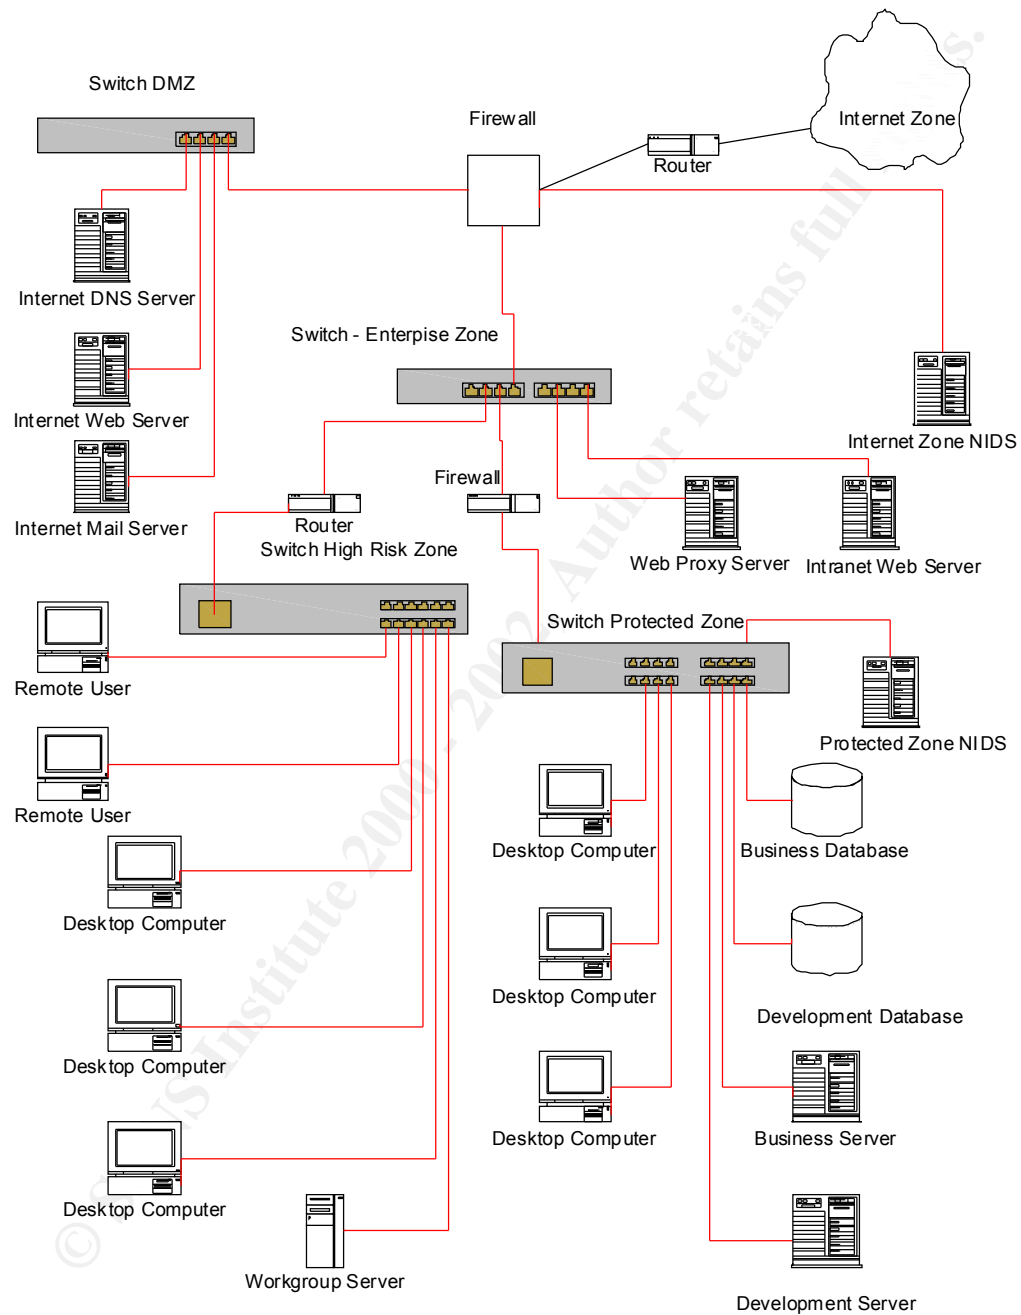
Using the same server as a mail server and a web server causes a "no-win" scenario when trying to respond to Internet attacks. Since the compromise of web servers has become common, these servers need to be thought of as disposable. If the same machine hosts the mail and web services, the web server cannot be taken offline without also losing email capabilities for the whole company.

Implement a Zone Defense

A good method for improving the topology of a network and adding an early response capability is to implement a *zone defense*. The basic steps for building a zone defense into a network are:

1. Identify and qualify your resources and assets based on risk and trust.
2. Place the resources in zones with other resources of similar risk and trust categories.
3. Place some sort of mechanism for enforcing rules between zones. These mechanisms are usually a combination of routers and firewalls.
4. Identify what traffic is necessary between the zones and configure the routers and firewalls to enforce these traffic requirements and deny everything else.
5. Implement a logging mechanism for traffic between zones.
6. Install Network Intrusion Detection Systems (NIDS) at key nodes on the network.

An example of a protected network topology which uses five zones is shown below.



The five zones are described in the following sections.

Internet Zone

The Internet Zone is the network encompassing the vast network outside of the corporate network. This is the source of outsider attacks so there is a firewall placed between this zone and the corporate zones.

Demilitarized Zone (DMZ)

This zone is for servers that need to have Internet exposure to perform their function. Examples of this would be Domain Name Servers (DNS), Internet Web Servers, Internet Mail servers and public FTP servers. Since these servers can potentially be compromised through their Internet exposure, the firewall protects the corporate zones through its firewall policy. Servers in the DMZ are protected from the Internet Zone through firewall policies which specify what ports are allowed to be forwarded from the firewall to DMZ servers, what protocols are allowed, what IP addresses are allowed to access services on the DMZ server, etc. If servers in the DMZ are infected by worms and viruses or compromised by hackers, the damage is contained in the DMZ.

Enterprise Zone

This is a corporate zone located behind the firewall where servers are placed that can be shared by the other corporate zones such as a web proxy server and a web server for internal use only. Notice also that a Network Intrusion Detection System is placed in this zone to monitor for threats.

High Risk Zone

This zone is used to segregate users, workstations and servers which might easily be infected or compromised. Possible reasons for placing a user or server in this zone might include:

- Exposure to outside networks by RAS clients
- Historical tendency for a user to become infected through unsafe computing practices
- The presence of a high number of unmonitored laptop computers
- Part time or contract technical workers with network access

Protected Zone

Key corporate assets are located in the Protected Zone in order to minimize their exposure to Internet threats and to problems in other zones. If the high risk zone becomes infected or compromised, one of the tasks for the response team to perform early in the attack would be to further restrict or eliminate network traffic to other zones by modifying the policy on the firewall that protects this zone. This example topology also uses a NIDS to monitor threats in this zone.

The implementation of these five zones and the layers of protection that the multiple firewalls and subnets provide is a good example of *defense in depth*. There can be failure or compromise in some zones, but key corporate assets are still protected.

Processes

Good technology and good network design can provide better protection, but the most important component of corporate asset protection is the implementation of business processes that specify actions to be taken during an Internet attack. *The most important process is to implement a security policy and update it often.* Having a security policy in place and having IT and corporate staff trained and briefed on the policy will reduce the confusion in the early stages of an Internet attack and catalyze the quick action required to eliminate or minimize the threat.

Build the Team

Identify individuals with the skills and motivation to be on an Internet Threat Response Team. Larger corporations may have the personnel and resources to commit to a dedicated response team while smaller companies will need to add these duties to the existing staff of system administrators and network engineers. Make sure the team members have this responsibility as part of their job description, performance reviews and compensation. Most important, train the team members in information security and incident response.

Establish Leadership

The security policy should clearly state who is in charge during the response to an Internet attack. Ideally this will be a corporate Security Officer and selected staff. For smaller companies this will be a group selected from the IT staff and trained for emergency response. By having clear leadership, the confusion during the critical early stages of an attack can be minimized.

History has shown that many attacks happen during non-business hours. Because of this, a system to contact the response staff must be clearly defined and readily available to IT and corporate staff. Team members need to be able to contact each other quickly and be able to communicate effectively. Names, phone numbers, pager numbers, and email addresses should be kept on easily distributed lists. A portion of the security policy should be dedicated to threat response and clearly describe who should be contacted and who is authorized to perform critical functions.

Authorize Containment

If a threat is discovered and it is determined that the network is vulnerable to this threat, it may be necessary to contain the threat to a certain zone of the corporate network while the threat response team determines how to proceed. This containment strategy can be very effective in reducing the damage to key corporate assets from a new Internet threat. An example of containment would be to block network traffic between a network zone that is infected with a new virus and the protected network zone. In this way, the key corporate assets located in the protected zone can be protected from a worm that travels from host to host through mechanisms such as unprotected shares. This quarantine action can be accomplished by modifying Access Control Lists [6] on routers or rulesets on firewalls [7] that have been placed between zones.

A more severe example would be to completely remove a machine from the network that has been compromised by a hacker. It may be necessary to surrender this compromised machine to law enforcement agents if they become involved in the incident response.

Insure Executive Understanding and Approval

Meet with company management regularly and keep them informed of the threats that are being directed at the corporate network. Make sure the executives understand the responses that will be taken during an attack and the implications of these actions. Get their approval in writing [8]. In order for the threat response team to be able to make the quick decisions and take the first actions that will reduce the damage, the team must be sure their actions are authorized and they will not be risking their jobs. Disabling Internet access or shutting down a corporate web site can be a very stressful task if the necessary steps have not been taken beforehand. Worms, viruses and hackers can do significant damage in the time it takes to get executive approval during the attack.

Train the Users to be Security Conscious

A good perspective to take is that everyone that uses a computer on the corporate network is part of the security. It only takes one user to open an infected email attachment or plug an infected laptop into the network to cause an event that could threaten key corporate assets. In order to address this risk, schedule security training for new employees and hold regular meetings or seminars for existing employees. Internet threats are receiving plenty of media attention, so most users are aware that there is a problem and do not want to be the one that causes a problem for the company. Keep the sessions brief and informative, but make sure participation is required. If there is a corporate Policies and Procedures Manual, make sure it contains guidelines for use of computer assets and Internet access.

Because the actions of employees while accessing the Internet has become such a problem, many companies are limiting or eliminating Internet access for certain job descriptions. While a very unpopular practice, it does reduce the risk to corporate assets from Internet threats.

Conclusion

The task of protecting corporate assets and Internet connectivity seems overwhelming to many corporations. This paper has shown that this monumental task can be accomplished through a combination of using key security technologies that are implemented on a well designed network with an organized security team in place. Corporations that choose to discount the size of the Internet threat and do not implement the types of measures described here *will* fall victim to an Internet attack. It is just a question of when, not if.

References

- [1] FedCIRC, “Defense Tactics for Distributed Denial of Service Attacks”,
<http://www.fedcirc.gov/docs/DDOS-defense.PDF> (2000)
- [2] CERT Coordination Center, “Trends in Denial of Service Attack Technology”,
http://www.cert.org/archive/pdf/DoS_trends.pdf (October 2001)
- [3] Riptech, Inc., “Information Security Challenges in the Electric Power Industry”,
<http://www.riptidech.com/custexp/whitepapers.html> (January 2001)
- [4] Riptech, Inc., “Understanding SCADA System Security Vulnerabilities”,
<http://www.riptidech.com/custexp/whitepapers.html> (January 2001)
- [5] SANS Institute, “SANS Security Essentials I: Information Security, The Big Picture”
(February 2002)
- [6] National Security Agency, “Router Security Configuration Guide”,
<http://nsa2.www.conxion.com/cisco/guides/cis-2.pdf>, (November 21, 2001)
- [7] National Institute of Standards and Technology, “Guidelines on Firewalls and
Firewall Policy”, <http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf>, (January
2002)
- [8] SANS Institute, “ SANS Security Essentials II: Network Security”, (February 2002)
- [9] Daniel F. Delong, “Under Pressure, Microsoft Moves To Tighten Security”,
<http://www.ecommercetimes.com/perl/story/13946.html> (October 4, 2001)
- [10] James Byne, “An Overview of Threat Risk and Assessment”,
<http://rr.sans.org/audit/overview.php> (January 22, 2002)