Network Hardening: Using Warfare Strategy

From Sun Tzu's The Art of War

Shawn W. Toderick

CCAI, RCHT, NET+

East Carolina University

ABSTRACT

There are a lot of Information Security and Network Security text and papers that quote some of the work of Sun Tzu's The Art Of War, mostly "Know your enemy and know yourself". Information security, while a part of business, is looked at not from a business context, but a warfare context. Information security is seen as a constant battle between the organization's administrators and hackers.

This paper attempts to apply some of the strategies Sun Tzu discusses to hardening an organization's network and information security. These strategies include knowing your network, knowing your enemy, misinformation and misdirection, perimeter security, and using the attacker's strengths against him.

"Know your enemy and know yourself; in a hundred battles, you will never be defeated. When you are ignorant of the enemy but know yourself, your chances of winning or losing are equal. If ignorant both of your enemy and of yourself, you are sure to be defeated in every battle."

–Sun Tzu; The Art Of War, Chapter 3

Within information security, Sun Tzu is often quoted because information security is often seen as a conflict between organizations and attackers where the battlefield is always the organization's network. Concepts like vulnerabilities, exploits, attackers, intruders, red teams, white hats and black hats are often used to convey this idea that security is a constant battle between good and evil within the realm of cyberspace. Sometimes, information security is compared to medicine where viruses infect individual systems within the organization with the intent of spreading to other systems to affect the entire body of the organization. Although this metaphorical comparison is usually contained to a specific area of information security such as anti-virus software and intrusion detection/prevention systems.

So which is it? Should we look at the protection of our piece of cyberspace as digital warfare or as digital medicine? And how does this comparison between life and death apply to a business environment? Since those areas of information security compared metaphorically to medicine are a subset of information security, digital medicine is most of the times after the fact, where as digital warfare is preparing the network for and engaging in battle so network security can be looked at with a mindset of war and not medicine. Using the principles of warfare, an organization's network can be hardened against almost all known attacks and is in a good position to defend against unknown attacks.

The argument between these two schools of thought is academic to the business world because it matters not how you treat the security of your digital information, but how well you protect it and the cost of that protection. If the cost of the protection is more than the cost of what is being protected, most organizations will not implement the protection. In business, information security boils down to simple Return on Investment.

At least that is how it used to be. In the world that we live in today, within the United States of America with identity theft and businesses holding so much personal information about its employees and customers, information security can no longer be looked at as a simple ROI situation. The impact of lost or stolen data can be catastrophic for a business. Privacy Rights Clearinghouse (www.privacyrights.org) lists a "Chronology of Data Breaches" reported, starting with the ChoicePoint breach on February 15th, 2005. [A Chronology of Data Breaches, 2005] 60 breaches were reported in a 6 month time frame affecting over 50 million people. Individual states are now considering security breach notification legislature [2005 Breach of Information Legislature] to identify when and how to report breaches.

A Security Panacea?

All this falls on individual organizations to improve their security posture and defenses (information security is almost solely defensive in nature). To do this, security and network administrators need an understanding of how intruders gain access to their networks to compromise the information. Administrators also need an understanding of how to deal with these breaches, both large and small. Using tactics borrowed from warfare is necessary; to fight intruders when they attempt to gain access to a network, to fight these intruders when they break through outer defenses and to fix the problems they cause. This philosophy of complete system integrity has been coined "Defense in Depth" by the U.S. Department of Defense and is addressed in the Department of Defense's Information Assurance Strategy.

Firewalls, while important to an organization's network security, are not the end-all be-all of network security. There is no such thing as a security panacea. Vendors implement the TCP/IP stack differently and as attackers advance in their understanding and manipulation of the TCP/IP stack and it various implementations requires vendors to release patches and upgrades to counter these emerging vulnerabilities. The time it takes for these new vulnerabilities to be recognized and fixes released leaves an organization open to be exploited through these vulnerabilities (sometimes referred to as Zero-Day Attacks). Something has to exist within the organization, not just at its network boundaries, to counter the intruders when this situation occurs. Organizations often implement Intrusion Detection and Prevention Systems (IDS and IPS) as one means to deal with intruders. Other methods include the use of anti-virus software to keep viruses, worms, and Trojans for affecting the network, and educating network users on how to recognize intruders and how to report them.

Know Yourself

Following Sun Tzu, traditionally knowing yourself requires knowing what organizational resources that need protecting. Performing risk analysis on business resources and identifying those resources that are most critical is the foundation of knowing one's self. A complete understanding of one's network is more than identifying what you have that needs protecting, it includes understanding how it can be attacked. In today's highly networked environment everything seems to have multiple communication paths to everything else. This gives intruders numerous paths of access, requiring administrators to evaluate each one to determine how much risk they pose to the organization, and how trust relationships between resources can be exploited when a low-risk resource has access to a high-risk resource.

For example, during risk analysis it is determined that a particular server is more critical than the workstations that access the server via a web interface. The server is patched, updated and host-based IDS is installed, but the workstations are left alone because it has been deemed to be cost ineffective to implement or upgrade current security on the workstations. All an attacker needs to do is compromise one of these workstations (which should be easy since security is lax on these hosts,) install a key-logger or other monitoring software and wait for someone with high-level privileges on the server to access it. Once this happens, the attacker now has indirect access to the server. Given time and the creativity of the attacker, the server itself may be compromised.

Since it is the implementation of the TCP/IP stack and the numerous protocols that exist within it that are exploited the most, aside from educating users, these outer and inner defenses requires a deep understanding of TCP/IP. This is evident in the granular control of ingress and egress traffic filtering that firewalls implement.

To truly know one's self, administrators need to perform penetration analysis (also referred to as pen testing.) In this situation, individuals legally attack a network trying to locate vulnerabilities and possibly exploit these vulnerabilities. The intent here is not to compromise the target network, but to perform the same actions attackers would perform to learn about the organization's weaknesses, using the same tools and programs attackers use. The results of a pen test would be then used to tighten network security by implementing new filters on firewalls, fine-tune IDS/IPS, increase user awareness of social attacks, and harden servers and workstations.

Hiring an outside organization to perform this pen test is full of legal pitfalls and ethical dilemmas. One of the first things that must be identified is the scope of the test. Is the entire security posture of the organization analyzed, or just part of it? Are vulnerabilities to be exploited to see the effect that it would have on the organization? Also, how well can you trust the individuals to keep the information the have learned about your organization confidential. For these reasons, and more, some organizations prefer to perform pen testing in-house. In-house pen testing also has a double whammy for Sun Tzu's theory of know yourself, know your enemy. When you pretend to be the enemy to learn what weaknesses you have, you also learn about your enemy.

Short TCP/IP Basics

Before delving into hardening the network, a brief discussion on the basics of TCP/IP is necessary. RFC 791 specifies the DoD Standard Internet Protocol and summarizes the IP header (in Section 3) and RFC 1180 provides a good TCP/IP tutorial. The IP header is 32-bits long:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version|  IHL  |Type of Service|          Total Length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Identification        |Flags|      Fragment Offset    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Time to Live |    Protocol   |         Header Checksum       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source Address                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Destination Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Options                    |    Padding    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

IP is an unreliable transport protocol used to carry all upper-level protocols and provides the transport and delivery of datagrams but does not provide a mechanism to verify that the datagram is from the host that sent it. This is the heart of IP spoofing: manipulating the source address field to hide the true identity of the source of the communication.

Certain fields in the IP header can be exploited based on their designed behavior. The TTL field is such a filed: by setting the TTL field to 1, an attacker can send an IP packet into an organization's network to learn of the first router, and increment the TTL field by 1 each time to gain more information about the organization's internal network. Reconnaissance such as this can be accomplished by using traceroute programs that serve legitimate purposes for network troubleshooting and analysis. Modifying the TTL field can also be used to deliver malicious content to a victim machine. By setting the TTL field to a value of one hop less than the location of the target the packet will not reach the destination, but may be seen by IDS in the path. The attacker can than increase the TTL field by one, and insert malicious content in the payload (data portion) of the packet. The IDS would see the new packet thinking it was a duplicate packet and that the host will ignore duplicate packets. Since the host never received the first packet, the second packet with the malicious content would be processed.

One of the most popular attacks discussed in classroom and water cooler environments is that of a SYN attack. SYN attacks take advantage of the TCP Three-Way Handshake process. When a client (source) requests services from another host (destination), the source initiates the handshake by sending a connection request to the destination. In this connection request, the TCP Flag is set to Session Initiation, or SYN, and the destination will allocate resources to this connection and responds with an acknowledgement of the request and sends its own initial sequence number (SYN/ACK) to the source. The source then responds with an ACK and the TCP connection has been established. In a SYN attack, the source never responds to the SYN/ACK leaving the destination to hold those resources open for a period of time. Multiple SYN requests can be made requiring the destination to reserve more resources for these connections. Since the attacker does not send the ACK of the SYN/ACK, the destination has fewer resources to respond to legitimate TCP requests. In some situations, the destination will completely exhaust all resources responding to these attack requests thereby denying all legitimate requests and some times crashing the destination host.

Covering all the protocol abuses that exist and the improper protocol implementations is beyond the scope of this paper. These examples should prove that an in-depth understanding of TCP/IP is fundamental to a solid understanding of network operations.

TCP Three-Way Handshake and Teardown

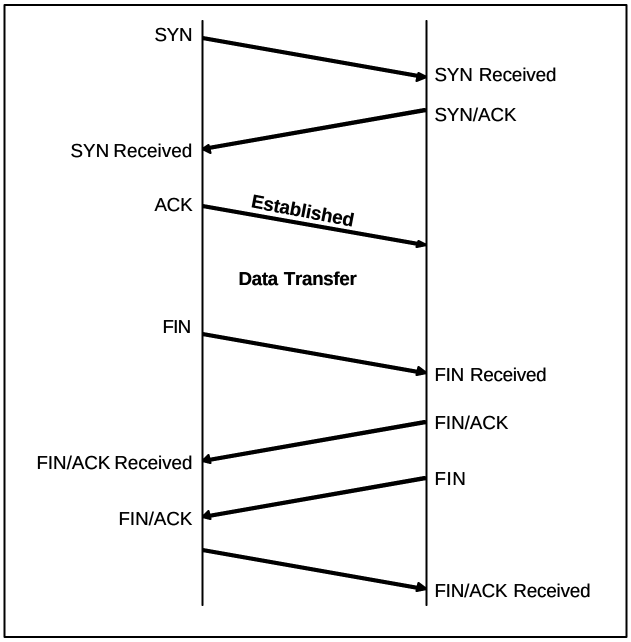

Gathering Information

"All warfare is based on deception."
Art of War, Chapter 1

There are many sources of information for an organization that an attacker can use to learn about the organization. Public records required by government agencies such as the Security and Exchange Commission (SEC) can give enough data that attackers can glean insight into a network. Corporate mergers and acquisitions leave both organizations open to attack when networks are merged.

One way to counter this is to reduce the amount of information that is publicly available. While all sensitive information cannot be withheld, some information can be. One military tactic that can be used here is to control the information that your enemy has about you. This can be done by misinformation or restricting access to certain information so that this sensitive information is seen by as few eyes as possible. Submitting public documents with incorrect information is illegal, so confer with the legal and financial departments of the organization to see what information is required to be made publicly available to ensure that only that information that is necessary is disclosed. The financial department is included because of full-disclosure and GAAP requirements may make control of what financial information is disclosed very difficult if not impossible. Attempting to control this information may be ignored from a network security standpoint due to legal constraints, thereby freeing up resources to concentrate on that information which can be controlled without legal ramifications.

In the book *Stealing The Network: How to Own a Continent*, [STN HOC] the character Sendai looks for a correlation between Secure Socket Layer (SSL) traffic at e-commerce sites and earnings. He does this by using the fragmentation ID field in the IP header to gather a sample of the amount of traffic the sites experience to estimate the total traffic per day and uses the organization's quarterly reports to find the correlation. While this book and the character are fictional, the tools and techniques used are accurate. This example does not, however, give the attacker information or access into the network but should point out that there are many sources of information that attackers can use to exploit something within the organization.

More information about an organization can be learned by searching ARIN (www.arin.net) for the IP address(es) the organization owns and uses for Internet access. A quick search can give all the IP addresses the target owns and which ones to attack.

While public document misinformation cannot be implemented, misinformation about network equipment can be. One of the most beneficial pieces of information that can be modified to confuse attackers is the operating system's (OS) fingerprint. Each OS implements standards differently. By sending multiple probes to a system and watching the responses, an attacker can learn which OS is being run on the system. By knowing what OS is being used, an attack can be customized for that OS.

One such tool that can do this is nmap by Fyodor (www.insecure.org). Nmap is a "free open source utility for network exploration or security auditing." By using the –O option (in newer versions of the program) nmap will perform remote host fingerprinting using TCP/IP. By using the –sV option, nmap can also attempt to determine the version of the services being used on the host. In the following example the IP address has been hidden (**xxx.xxx.xxx.xxx**) and some information removed to protect the target:

```
# nmap -O –sV xxx.xxx.xxx.xxx

Starting nmap 3.81 ( http://www.insecure.org/nmap/ )
Interesting ports on xxx.xxx.xxx.xxx:
(The 1657 ports scanned but not shown below are in state: closed)
PORT     STATE SERVICE
22/tcp   open  ssh         OpenSSH 3.9p1 (protocol 1.99)
80/tcp   open  http        Apache httpd 2.0.52 ((Fedora))
111/tcp  open  rpcbind     2 (rpc #100000)
5801/tcp open  vnc-http-1?
5901/tcp open  vnc         VNC (protocol 3.8)
6001/tcp open  X11         (access denied)
[... OUTPUT OMITTED ...]
Device type: general purpose
Running: Linux 2.4.X|2.5.X|2.6.X
OS details: Linux 2.5.25 - 2.6.3 or Gentoo 1.2 Linux 2.4.19 rc1-rc7)
Uptime 42.865 days (since Wed Jun  8 15:03:48 2005)

Nmap finished: 1 IP address (1 host up) scanned in 72.289 seconds
```

Nmap guesses at the OS and version (`Running: Linux 2.4.X|2.5.X|2.6.X` and `OS details: Linux 2.5.25 - 2.6.3`) of the target, with OpenSSH 3.9p1 and Apache 2.0.52 services running as well. (The target is actually Fedora Core 3 with a 2.6.11 kernel.) An attacker can now focus the attack against this machine using this information looking for exploits of the services and service version the host is using.

This OS fingerprinting can be scrambled by either modifying the OS itself (which may be illegal depending on the licensing of the OS) or at the perimeter with the firewall. One firewall vendor, Check Point, has a feature called SmartDefense that can scramble some of the IP header fields commonly used for OS fingerprinting. Another feature of Check Point's SmartDefense is to perform ISN (initial sequence number) scrambling. [CPSD, pg 172 -173] Each OS uses an algorithm to generate the ISN that the host will use in a TCP connection (the Three-Way-Handshake). Unfortunately, the algorithm used to generate the ISN is highly predictable, so an attacker can send successive SYN requests and compare the differences in the ISN responses to guess the OS based on the algorithm. The ISN scrambler can counter this type of information gathering by creating a different ISN when translating the packet for outside communication.

It is not possible to completely prevent this type of information gathering, but we can make it more difficult for attackers and hope they move to an easier target. Unfortunately for some firewalls that cannot perform this OS masking, attackers can learn about the firewall itself:

```
# nmap -O -sV xxx.xxx.xxx.xxx

Starting nmap 3.81 ( http://www.insecure.org/nmap/ )
Interesting ports on xxx.xxx.xxx.xxx:
(The 1661 ports scanned but not shown below are in state: closed)
PORT    STATE SERVICE VERSION
22/tcp  open  ssh     Cisco SSH 1.25 (protocol 1.5)
443/tcp open  ssl     OpenSSL
Device type: firewall
Running: Cisco PIX 5.X|6.X
OS details: Cisco PIX Firewall (PixOS 5.2 - 6.1), Cisco PIX Firewall
running PIX 6.2 - 6.3.3

Nmap finished: 1 IP address (1 host up) scanned in 9.152 seconds
```

This scan of the firewall accurately identifies the vendor and OS version of the firewall and what ports can be used to gain access to it. Fyodor discusses remote OS detection in his paper "*Remote OS detection via TCP/IP Stack FingerPrinting*" (http://www.insecure.org/nmap/nmap-fingerprinting-article.html).

Other sources of information are not as easily controlled. CERT (Computer Emergency and Response Team) Advisories (http://www.cert.org/advisories/) about the latest vulnerabilities and exploits, which are freely and publicly disclosed, can give an attacker a way in if the organization has not addressed these issues. Full-disclosure mailing lists such as Bugtraq provide a wealth of knowledge about cutting edge vulnerabilities.

Internet storm centers such as the SANS (System Administration, Networking, and Security) ICS provide global data on the current state of health of the Internet and the most widely implement attacks. Certain software can be used to dynamically update firewalls and IDS/IPS with information from ICS. Check Point's SmartDefense, with the appropriate licensing, can send log entries to DSHIELD, "an attack correlation engine with world wide coverage." (secure.dshield.org) and receive dynamic updates to the Rule Base. Linux firewalls can be scripted to automatically send log entries to DSHIELD and receive reports and automatically update firewall rules.

All this information is available for attackers to find the latest weaknesses within your organization. Administrators rarely have the time to devote to discovering all the weaknesses that their systems have, so keeping up with these lists is a must to ensure that the network is as well protected as possible.

Honeypots and Honeynets are other tools that can be useful in misdirecting attackers into an area that protects the organization without compromising it. A honeypot is a "trap set to detect or deflect attempts at unauthorized use of information systems. Generally it consists of a computer, data or a network site that appears to be part of a network but which is actually isolated and protected, and which seems to contain information that would be of value to attackers." [Honeypot, 2005] "Two or more honeypots on a network form a honeynet. Typically, a honeynet is used for monitoring a larger and/or more diverese network in which one honeypot may not be sufficent." [Honeynet, 2005] The issue with these distraction tools is that the attacker must first penetrate the networks outer defenses.

Honeypots/nets can serve another purpose for network security if it is close enough in similarity to the actual network. Vulnerabilities and weaknesses can be discovered in the honeypot/net that can be directly transferred to the actual network without compromising the real network.

More information about Honeypots and Honeynets can be found at www.honeypots.net and various Honeypot/net projects such as http://www.lucidic.net/, http://zaro.alestra.net.mx/honeynet.html, and http://www.honeynet.org/. When implementing these distraction measures, be aware of the legal issues with them. In some situations, using these tools can be considered entrapment, and administrators still may not be able to electronically pursue the attackers. Always consult with legal counsel prior to tracing and tracking attackers electronically.

Misdirection can also wear an attacker down to the point that he will give up and seek another target. The best weapon that attackers have is that they are invisible until they strike. Once they strike, time is against them. The longer it takes to accomplish his goal, the greater the chance of being detected. Bad network maps, wrong OS detection and honeypots/nets take time to discover their existence. Honeypots/nets are almost always a sign of detection. This wasted time may be too much for the attacker and he will leave. Wasted time can also aggravate and upset the attacker. Most attackers have a healthy dose of fear of being caught that keeps them cautious. When this fear has lessened, mistakes are usually made and mistakes lead to detection.

Starting at the Perimeter

"Generally, he who occupies the field of battle first and awaits his enemy is at ease, and he who comes later to the scene and rushes into the fight is weary"
Art of War, Chapter 6

One advantage security administrators have over attackers is that they are more familiar with the battlefield (the organization's network) than the attacker. The attacker has to learn about the network initially through secondhand information, and if this information is outdated or incorrect (as through misinformation) the attacker is at a severe disadvantage.

This advantage of ownership of the battlefield can allow administrators to prepare for the attack, which is most of the time a surprise attack (unless Internet Storm Centers are used to give a heads up on possible attacks.) The severity of surprise attacks can be lessened by the use of perimeter security, i.e. firewalls, since there is no such thing as true real-time alerts for an attack. All alerts are after the fact, even if the alert is generated the moment after the attack occurs and this time difference can be all an attacker needs to slip inside and dig in.

Firewalls serve four purposes: provide a barrier between the organization's network and the Internet or another open un-trusted network, filter traffic, act as a first alert system to attacks and are the first line of defense against attacks. By identifying where the organization's network begins, it is easier to clarify where the organizations domain of control extends to. This is of utmost importance when engaging in a track and trace of an attacker. In some instances an organization does not have legal permission to pursue an attacker outside its network boundaries. So any arrests made as a result of an illegal track and trace will be void and the attacker can slip back into the animinity the Internet offers.

Firewalls also serve as a filter for traffic both ingress and egress. In the event of an attack originating from outside the network, the firewall can filter out that traffic associated with the attack while allowing legitimate traffic to pass through. The only alternatives here are to either allow the attack traffic into the internal network and deal with it there or pull the organization's Internet connection (which would be a very difficult action to justify regardless of the severity of the attack.) If the organization's systems have been compromised and are part of a Distributed Denial of Service (DDoS) attack, the firewall can again filter out this traffic originating from within the organization from reaching the Internet.

When properly configured and deployed, firewalls can also generate first alerts to the possibility of an actual attack, since the firewall is the border device and all traffic must traverse the firewall to inter/exit the organization's network. Attackers must first by-pass or compromise the firewall before he can gain access to the internal network. Through the use of logs, attackers who prefer slow, hard to detect attacks can be detected. Although this requires a lot of man hours, or cpu cycles, to peruse through thousands of lines of log entires.

Firewalls are the first line of defense against attacks since, as pointed out previously, all traffic must traverse the firewall. Certain precautions can be implemented on the firewall to help protect the internal network from initial attacks. Some firewalls are capable of monitoring the queues for TCP communications, protecting against SYN Flood attacks. The Cisco PIX Security

Appliance has the capability to monitor half-open TCP connections (a.k.a. “embryonic” connections) to hosts on the inside of the network and can temporarily sever communication between a internal host and the initiating external host if these half-open connections exceed a certain limit or time frame. Half-open TCP connections are those sessions where the SYN and SYN/ACK have been sent, but the session has not been established, as described earlier in a SYN attack.

Some organizations are doing away with firewalls. Stuart Berman briefly describes how his organization has done away with their firewall in his article *“The Death Of A Firewall”* from Network Magazine. [Berman, 2005] This is both interesting and frightening. It is interesting in that the organization can accomplish its security requirements, and manage Internet usage, without filtering traffic at the perimeter. It is frightening in that there exists no barrier between the organization and the rest of the world. This paper does not take into consideration the absence of perimeter security for these reasons. But it is interesting to note that there exists within network security the idea, and application, that perimeter security is not necessary.

The firewall is an important part of the overall network security and security posture of the organization, but cannot be the sole defense against attacks. Firewalls are difficult to configure and understand, making complex configurations difficult to implement. Organizations need more than just perimeter defenses.

Once Inside

"A speedy victory is the main object in war. If this is long in coming, weapons are blunted and morale depressed.”
Art of War, Chapter 2

Once an intruder has passed the perimeter security, the firewall can no longer protect the network. We need tools that exist within the network to catch these intruders. Tools like IDS/IPS, ant-virus software, patched and updated workstations and servers, and help desks go a long way for network security at this point.

Also, if misinformation is used, once the attacker has passed through the perimeter he may spend time trying to learn the real design and setup of the internal network. The problem here is that the longer an attacker is inside the network, the greater the chance of being caught. Without the use of IDS, the longer the attacker is inside the network the greater the chance the attacker’s traffic will be seen as legitimate.

Intrusion Detection Systems are the “tools, methods, and resources to help identify, assess, and report unauthorized or unapproved network activity,” and have been around for a while. [IDS/IPS]. To clarify discussion, an intrusion is “an active sequence of related events that deliberately try to cause harm.” [IDS/IPS] Intrusion Prevention Systems are different from IDS in that they sit inline on the network, meaning they sit in the middle of the data path, and can perform predefined actions on traffic. IPS is still relatively new, being less than 7 years old as of this writing.

IDS should not be confused with IPS. While some IDS can allow you to perform some rudimentary actions on traffic, such as dropping a packet or resetting a TCP session, this is not prevention. These actions that IDS take are still after the fact so it is not truly prevention. There are also two forms of each type: host-based which is software that resides on workstations and servers, and network-based which provide security for the network communications.

Host-based IDS/IPS (HIDS/HIPS) monitor log files and system changes on the host on which they are installed looking for misuse or intrusion. Some such programs include Tripwire (www.tripwire.com), GFI LANguard S.E.L.M (www.gfi.com/lanselm/?adclickid=3737283), Linux IDS (LIDS, www.lids.org), and Dragon Squire (for UNIX). These programs can actively monitor the host system looking for misuse and quickly alert administrators of the misuse. While these programs do not perform in real-time, they do detect misuse quickly enough that administrators can remove them before they have a chance to affect other host systems.

One issue with HIDS/HIPS is that they only protect that individual host; they cannot affect the traffic entering the host or traffic on the wire but can halt attack traffic leaving the system. Without a central monitoring agent, compromised systems would have to be manually checked to see if patterns emerge that indicate that an organization wide attack is in place.

Network-based IDS/IPS (NIDS/NIPS) monitor network traffic on the wire and can affect the traffic before it reaches the host, or consume more bandwidth than it already has. By monitoring the packets crossing the network, unsuccessful attacks can be detected (unlike HIDS/HIPS which only verify the success/failure of an attack). NIDS/NIPS are closer to real-time responses than HIDS/HIPS but some delay exists.

During the reconnaissance portion of an attack, or during the actual attack, the intruder's traffic can be quickly detected by NIDS/NIPS if common attack vectors are used. Zero-Day attacks, unusual attack vectors are difficult to detect and improperly tuned NIDS/NIPS can so many register false positives as to render the defense weapon unusable. These tools do not come pre-tuned out-of-the-box, so they must be configured to ignore the organization's normal traffic patterns that it might detect as misuse or unusual.

Anti-virus software goes a long way to protect individual hosts and stop the spread of a virus/worm/Trojan infection in the organization's network. Some AV software can even detect if a host has been turned into a zombie for a DDoS attack. AV software is perhaps one of the most useful tools from the end-user's perspective because it is one of the few security software packages that they are the most familiar with and understand. End users with computers at home are usually savvy enough to understand that they need this level of protection and when they use it both at home and work, they feel more comfortable with the technology. When an alert is generated with AV, they understand that this is serious and need to alert the organization's Help Desk/Tech Support department.

End users are an excellent security tool. They are the actual users of the network and end systems. When all other security defenses fail, it is usually the end user that notices something is wrong and report to the Help Desk. End users are also the victims of social engineering attacks, so it is imperative that they are aware of their surroundings, who they are talking to, and what to do if they feel that a security event has arisen. Most of the time training is used to accomplish this, but organizational policy and procedures also help the end user understand their role in the overall security of the organization.

Preparing for an Attack

"Generally, management of a large force is the same as management of a few men. It is a matter of organization."
Art of War, Chapter 5

It is not a matter of *if* an organization will be attacked, but *when*. There should be no disillusions that the organization is 100% completely safe from intruders, regardless of the level and strength of defenses the organization has implemented. All these precautions are to lessen the severity of the attack and to alert that an attack is or has taken place so that the organization can properly handle the situation.

To this end, policies and procedures are paramount. These documents dictate the organization's security posture and how these defenses can be implemented. The Acceptable Use Policy (AUP) dictates what end users can and cannot do; how to report suspicious activity or possible security event; the domain of control the organization maintains.

The Incident Response Procedure (IRP) is the implementation of the organization's response policy. The IRP should, among other things, indicate how to handle an attack that is ongoing as well as after the attack has subsided. This document should state what situations, if any, require pulling the plug on the network in the event of an attack. From a business perspective some organizations find this result to be unacceptable, and from a security perspective in some situations may make matters worse if the attacker has thoroughly infected the network. The IRP should also indicate the positions, never the individual's name, that are part of the Computer Incident Response Team (CIRT); the notification procedure, and how the information about the response should be recorded and handled. Not following the proper chain of custody will render any evidence gathered inadmissible in a court of law. The CIRT should regularly practice incident response so that when an incident occurs they are prepared for it and will not let the excitement of the situation overwhelm them and their ability to make rational decisions.

Without these documents, and other Business Continuity documents, the organization cannot effective implement and enforce a comprehensive and coherent information security plan.

If penetration testing is permissible, either in-house or out-sourced, than it should be done regularly but not on routine schedules. If these tests are performed on a predictable schedule then attackers can design their attacks to coincide with the testing in an attempt to mask their tracks.

“Therefore, analyze the enemy's plans so that you will know his shortcomings as strong points. Agitate him in order to ascertain the pattern of his movement. Lure him out to reveal his dispositions and ascertain his position.”
Art of War, Chapter 6

When an attack occurs, if it is not permitted to pull the plug, then the CIRT and security administrators should monitor what the attacker is doing. This monitoring will help the organization better prepare itself and its defense for the next attack. Use honeypots/nets and pen testing to become familiar with the attackers, their tools, and methods of attack. All of these defenses blunt the attacker’s weapons of choice in one fashion or another and can make him weary of the battle.

Using the results from pen tests and data collected from honeypots/nets to fine-tune network parameters such as IDS/IPS sensitivity, scanning times and frequency of AV, and firewall filters. Set TCP connection queues to smaller time frames and increase the number of queues to combat SYN attacks while still allow legitimate connections. Evaluate the use of VPNs to provide a secure communications channel across the Internet to the organization’s network instead of using unsecured protocols such as TELNET and FTP. Evaluate the end user’s responses to determine if more training is necessary or procedures need to be modified.

The principle here is to use strategies from warfare to evaluate the organization’s information and network security to better prepare the organization for the next battle against its enemies. Prepare, test, evaluate, change, repeat.

References:

Honeypot. (2005). Retrieved Jul. 19, 2005, from Wikipedia, the free encyclopedia. Web site: http://en.wikipedia.org/wiki/Honeypots.

Honeynet. (2005). Retrieved Jul. 19, 2005, from Wikipedia, the free encyclopedia. Web site: http://en.wikipedia.org/wiki/Honeynet.

Berman, S. (2005, Jun 01). The death of a firewall. Network Magazine, Retrieved Jul 29, 2005, from http://www.networkmagazine.com/shared/article/showArticle.jhtml?articleId=163700676.

A chronology of data breaches. (2005). Retrieved Jul. 16, 2005, from Privacy Rights Clearinghouse Web site: http://www.privacyrights.org/ar/ChronDataBreaches.htm.

2005 breach of information legislature. (2005). Retrieved Jul. 16, 2005, from National Conference of State Legislatures Web site: http://www.ncsl.org/programs/lis/CIP/priv/breach.htm.

Other Resources

The quotes from Sun Tzu's Art of War used in this paper come from the Qing Long Institute's website:

http://qing_long_institute.tripod.com/qinglonginstitute/id11.html

DoD Directive Information Assurance (IA), number 8500.1. (2004). Retrieved Jul. 17, 2005, from Defense Tactical Information Center Web site: http://www.dtic.mil/whs/directives/corres/pdf2/d85001p.pdf.

Curtin, M. (1997). Introduction to network security. Retrieved Jul. 15, 2005, from Interhack Research Web site: http://www.interhack.net/pubs/network-security/.

Endorf, C., Schultz, E., & Mellander, J. (2004). Intrusion Detection & Prevention. Emeryville, Cal: McGaw-Hill/Osborne.

Mandia, K., Prosise, C., & Pepe, M. (2003). Incident response & computer forensics. 2nd ed. Emeryville, Cal: McGaw-Hill/Osborne.

McClure, S., Scambray, J., & Kurtz, G. (2003). Hacking Exposed Network Security Secrets & Solutions. 4th ed. Berkeley, Cal: McGaw-Hill/Osborne.

Harris, S., Harper, A., Eagle, C., Ness, J., & Lester, M. (2005). Gray Hat Hacking The Ethical Hacker's Handbook. Emeryville, Cal: McGaw-Hill/Osborne.

Whitman, M., Mattord, H.,(2003). Principles of information security. , Canada: Course Technology.