Research Paper: Information Security Technologies

by

Benjamin Tomhave

November 10, 2004

Prepared for:

Professor Dave Carothers
EMSE 218
The George Washington University

This paper or presentation is my own work. Any assistance I received in its preparation is acknowledged within the paper or presentation, in accordance with academic practice. If I used data, ideas, words, diagrams, pictures, or other information from any source, I have cited the sources fully and completely in footnotes and bibliography entries. This includes sources which I have quoted or paraphrased. Furthermore, I certify that this paper or presentation was prepared by me specifically for this class and has not been submitted, in whole or in part, to any other class in this University or elsewhere, or used for any purpose other than satisfying the requirements of this class, except that I am allowed to submit the paper or presentation to a professional publication, peer reviewed journal, or professional conference. In adding my name following the word 'Signature', I intend that this certification will have the same authority and authenticity as a document executed with my hand-written signature.

Signature _____Benjamin L. Tomhave_____

Research Paper: Information Security Technologies

by

Benjamin L. Tomhave

Abstract

The following research paper provides analysis of thirteen (13) information security technology topics, arranged in ten (10) groups, that are either commonly found or emerging within the information security industry. These topics include: Access Control Management, Antivirus, Audit Data Reduction, Firewalls, Intrusion Detection Systems (IDS), Intrusion Prevention Systems (IPS), Anomaly Detection Systems (ADS), Event Correlation Systems (ECS), Network Mapping, Password Cracking, Public Key Infrastructure, Virtual Private Network, and Vulnerability Scanning Systems. IDS, IPS, ADS and ECS are grouped together under one common heading (Intrusion Detection and Analysis Systems) due to their commonality and interdependence. This paper provides basic overview information about each technology, but primarily focuses on analyzing each technology within the modern information security and business context, looking at how it meets business needs while addressing Confidentiality, Integrity and Availability as a Countermeasure that Detects, Corrects and/or Protects.

Table of Contents

Research Paper: Information Security Technologies

by

Benjamin L. Tomhave

I.INTRODUCTION AND OVERVIEW OF APPROACH

This research paper introduces and analyzes ten (10) information security technologies. Each of the following sections focuses on a specific technology and adheres to the following general format:

- Technology Overview: A high-level introduction to the technology.
- Business Analysis: An evaluation of the usefulness, cost, complexity, and utility of the technology in the modern business environment.
- Security Analysis: The security technology is weighed against the tenets of Confidentiality, Integrity and Availability as well as evaluating its role as a countermeasure (detect, correct, protect).

The ten security technologies addressed in this paper are:

1. Access Control Management
2. Antivirus
3. Audit Data Reduction
4. Firewalls
5. Intrusion Detection and Analysis Systems
6. Network Mapping

7. Password Cracking
8. Public Key Infrastructure
9. Virtual Private Networks
10. Vulnerability Scanning Systems

II.ACCESS CONTROL MANAGEMENT

Access control management (ACM) systems pull together identity, authentication and authorization to restrict what resources a user may access and in what manner that access may occur (read, write, execute, modify, etc.). ACM solutions may be based on a number of security models, including Discretionary Access Control (DAC), Mandatory Access Control (MAC), and Role-Based Access Control (RBAC). A standard ACM provides an interface through which a user will self-identify, followed by a mechanism for challenging and confirming that identity, and then a method for granting rights, or access to information, based on the non-repudiated authentication of the user. Access control is at the heart of information security and is the fundamental premise upon which the industry is based[1]. Without access control management, there would no method through which to provide security for systems and data.

A.Business Analysis

Access control management systems provide the foundation for information security within the business environment. Its usefulness is extensive, with the primary functions

[1] Ben Rotchke, *Access Control Systems & Methodology* (New York: SecurityDocs.com, 2004, accessed 06 November 2004); available from http://www.securitydocs.com/go/69; Internet.

being to classify data systems according to value and allocate protection mechanisms in accordance with the value of the resource. According to Tipton and Krause, "[the] essence of access control is that permissions are assigned to individuals or system objects, which are authorized to access specific resources."[2]

The implementation of ACM systems can range in cost from minor to extreme, depending on the value of the resource being protected. The underlying security model applied also impacts how expensive and complex the solution may be. ACM solutions are perhaps the most important security technology that can be deployed, ahead of all other countermeasures, because of its inherent purpose to control access to data and systems. The utility of the ACM systems, however, is limitless under the assumption that a business has resources of value that require protecting.

Discretionary Access Control systems are very common and are generally cost-effective for most environments. Most operating systems today - ranging from Windows to UNIX to Linux and beyond - make use of a DAC model of access control. Mandatory Access Control systems tend to be more complex and costly in performance and maintenance. MAC systems require a much stronger systematic adherence to the precepts of access control and can thus challenge administrative resources and confound access to data as required by the business. Implementation of MAC requires proper foresight and planning to avoid difficulties in the long term; an effort that is often a costly engineering effort frowned upon by the business. Finally, Role-Based Access Control systems are

[2] Harold F. Tipton and Micki Krause, <I>Information Security Management Handbook, 4th Edition<I> (Boca Raton: Auerbach, 2000), p1.

increasing in popularity and are predicted to saving companies millions of dollars in the coming years.[3]

B.Security Analysis

An access control management system has the potential for impacting all three tenets of information security (Confidentiality, Integrity and Availability). The primary role of an ACM solution is to protect the confidentiality of a resource by restricting access to the resource. Additionally, an ACM solution will control the attributes of the access, such as read, write and execute. For example, in the case of a data file, an ACM system may grant a user read access, but deny access to write or modify the data within the file.

Under a DAC model, access controls are managed directly by the resource owner. In a MAC model, the system dictates what level of access may be granted to a resource. Finally, RBAC assigns access based on the rights of a group (or role) within the system. All users who share a given role have the same access. This approach contrasts to DAC where each user may have a unique set of rights. MAC is similar to RBAC in terms of using a role-based approached based on labeling. However, the inner operations of a MAC vary distinctly from an RBAC; discussion of which exceeds the scope of this document.

Access control management systems hinge on the proper identification of subjects trying to access objects. The process of positively identifying a subject is called authentication.

[3] National Institute of Standards and Technology, <I>NIST Planning Report 02-1: Economic Impact Assessment of NIST's Role-Based Access Control (RBAC) Program<I> (Washington: NIST, 2002, accessed 12 October 2004); available from http://csrc.nist.gov/rbac/rbac-impact-summary.doc; Internet.

The authentication process usually occurs when a subject self-identifies and then responds to a systematic challenge of the identity. This challenge is based on what you know, what you have or who you are. A password is an example of something that you may know, and is currently the most common method of proving identity. A token is an example of something that you have, and biometrics is an example of who you are. Biometrics is a method of identification based on the physical characteristics of a human being, such as a fingerprint, iris scan or retinal scan. Biometrics, though holding significant promise as part of an access control management system, also has significant drawbacks, such as to acceptability to users, reliability and resistance to counterfeiting.[4]

The future of access control management systems appears to be in the direction of multi-factor authentication, oftentimes making use of passwords in combination with tokens or biometrics. Beyond the current trend, it seems likely that passwords will eventually be rendered completely obsolete in favor of some form of token or biometric becoming the first, if not only, form of authentication. Specifically, use of numeric or data tokens is on the increase and projected to continue gaining in popularity and acceptance. Major international Internet Service Provider America Online has recently announced the availability of numeric tokens for users as a second factor for authentication. Additionally, as public key infrastructure solutions (see Section IX below) mature and gain in prevalence, the use of data tokens will increase in importance. For example, a bank will be able to issue a USB-based data token to a customer. On the data token will be the customer's unique identifier in the form of a digital certificate. This certificate will

[4] Donald R. Richards, "Biometric Identification," in <I>Information Security Management Handbook, 4th Edition<I>, ed. Harold F. Tipton and Micki Krause (Boca Raton: Auerbach, 2000), p9.

be managed through a central Certificate Authority and will be used both for authentication and for encrypting and digitally signing communication and transactions.

Thus, access control management will not only continue its central role within information security, but it will also grow in scope, adding more extensive capabilities for positively impacting confidentiality and integrity. Additionally, besides protecting resources, it may also include extended capabilities that will allow for easier detection of attacks and possibly even automatic methods for correcting violations of integrity.

III.ANTIVIRUS

The first computer virus credited with being found "in the wild" is believed to be a program called "Elk Cloner" that targeted Apple DOS 3.3.[5] The term "virus" may actually have originated in the 1970s in science fiction literature[6], though as a concept it has likely been around since the 1960s. Traditionally, "[a] virus is simply a computer program that is intentionally written to attach itself to other programs or disk boot sectors and replicate whenever those programs are executed or those infected disks are accessed."[7] In the modern context, this traditional form of malicious code, or malware, is less common. Instead, it is far more common to see variations on this original theme in the form of "worms" and "Trojan horses" that infect a computer system either through direct execution or through some form of network-based replication method. In the

[5] Wikipedia, *Computer virus* (St. Petersburg: Wikipedia, 2004, accessed 06 November 2004); available from http://en.wikipedia.org/wiki/Computer_virus; Internet.
[6] Wikipedia, *Computer virus* (St. Petersburg: Wikipedia, 2004, accessed 06 November 2004); available from http://en.wikipedia.org/wiki/Computer_virus; Internet.
[7] Bob Kanish, *An Overview of Computer Viruses and Antivirus Software* (Unknown: Kanish, 1996, accessed 12 October 2004); available from http://www.hicom.net/~oedipus/virus32.html; Internet.

modern context, hybrid malware programs typically replicate through worm-like behaviour that preys on vulnerabilities in operating systems or through social engineering attacks, and then setup backdoors via the Trojan horse mechanism. This backdoor can then allow the attacker to remotely access and control an infected system, allowing for the perpetration of other illicit activities, such as sending SPAM or using the compromised system as a proxy, or relay, through which remote access can be gained to otherwise-protected resources.

Antivirus software has been around for at least the past 10-15 years, though no references were found that indicated a specific date when such programs were first made available. Antivirus software was developed to detect the presence, and eventually the attempted infection, of a system by malware. There are generally two types of antivirus scanning software: signature-based and heuristic. Signature-based scanning relies on a database of known malware signatures. It must be updated on a regular basis in order to ensure a current database of known malware. According to eBCVG, an IT Security company, a heuristic scanner "looks at characteristics of a file, such as size or architecture, as well as behaviors of its code to determine the likelihood of an infection."[8] The downside to heuristic scanners is that they often generate results that misidentify software as being malware (a.k.a. "false positives").

The most popular operating system, in terms of pure numbers, is Microsoft Windows. As such, it is also the most targeted platform by malware. There are several companies who provide AV software for Windows. There are also versions of AV software for other

[8] eBCVG IT Security, *Heuristic Scanning - Where to Next?* (Tel-Aviv: eBCVG, 2004, accessed 12 October 2004); available from http://www.ebcvg.com/articles.php?id=264; Internet.

platforms, like Mac OS, UNIX and Linux. However, there are very few cases of malware for those platforms, due in part to their distinct differences from Windows.

A.Business Analysis

In the modern age of computing, antivirus (AV) software is very inexpensive, very common, generally easy to deploy, and oftentimes relatively easy to maintain (easier than patching operating systems and applications, but still more challenging than being fully self-contained). Furthermore, the prevalence and availability of antivirus as a very basic countermeasure is such that a legal argument could be successfully made that the failure of a business to implement AV software throughout the organization could be deemed an act of negligence. As such, the utility and usefulness of AV software is very obvious, both from the standpoint of minimizing the threat of malware and from limiting legal liability resulting from a malware infection.

AV software itself is generally not complex. Most AV packages rely primarily on signature-based scanning with minor heuristic scanning capabilities integrated. The software is usually simple to install and is configured by default to automatically update the underlying scanning engine and the signature database on a regular basis from the Internet.

B.Security Analysis

Whereas businesses are expected to install and maintain antivirus software on most, if not all, systems as a matter of limiting legal liability, the effectiveness of AV software

diminishes each day. The AV industry has generally reached a plateau in the last five years and has not made any major advances in the ability to detect and prevent malware infection. Furthermore, the growth in popularity of the Internet has caused the computing world to become highly interconnected, leading to the development of so-called "zero-day exploits." These exploits correspond to vulnerabilities that are released on the same day in which the exploit itself is released. In the worst-case scenario, a major organization like Microsoft will announce the presence of a vulnerability in their popular Windows operating system mid-day, and by that evening a worm will be circulating on the Internet that is actively looking for vulnerable systems and attempting to infect them through this new vulnerability. Sadly, such events have happened in recent history, and oftentimes before a patch is even available to fix the vulnerability and before AV signatures have been developed and released.

The purpose of AV is to detect, protect and correct. Specifically, antivirus software is designed to detect malware infections, but it is also able to protect against an active infection attempt, and it is also often able to correct by disinfecting a system, depending on the characteristics of the malware. From the standpoint of Confidentiality, Integrity and Availability, AV software primarily addresses Integrity. The goal of AV software is to protect the Integrity of the operating system, application or data. Additionally, it has a secondary benefit of ensuring the availability of an object by detecting, protecting or correcting malware infections. Confidentiality may also be protected indirectly for malware that may cause data to be sent out randomly, such as Word documents as attachments, forwarding emails, etc.

IV.AUDIT DATA REDUCTION[9]

Audit Data Reduction is an emerging field of study in information security. The Audit Data Reduction Group, part of the COAST Laboratory at Purdue University in the Center for Education and Research in Information Assurance and Security (CERIAS), appears to be a leader in innovative research and thinking on the subject. The problem being addressed relates to the amount of audit data created, out of necessity, by critical systems. These critical systems often generate copious amounts of audit logs, which are often difficult to pour through for signs of malfeasance. The goals of audit data reduction systems are to contribute to misuse and anomaly detection. These types of systems are discussed further in Section VI.

A.Business Analysis

Audit data reduction (ADR) will increasingly become a useful and necessary part of the information security solution toolset. Businesses are increasingly inundated with audit logs generated by all critical systems. The advent of federal regulations that require thorough logging, such as within "financially significant systems," will further contribute to this trend. As a result, in order to maximize the value of these audit logs with an eye toward reducing risk to the overall business, it will become increasingly necessary to condense these raw logs into a more useful format.

[9] Purdue University, *CERIAS: Audit Trail Reduction Group* (West Lafayette: CERIAS, undated, accessed 12 October 2004); available from http://www.cerias.purdue.edu/about/history/coast/projects/audit-trails-reduce.php?output=printable; Internet.

Today, audit data reduction systems are still early in academic and commercial development. Solutions tend to be relatively complex and costly. However, it seems very likely that these systems will improve over time and decrease in complexity. In the end, we will likely see large audit data repositories built, based on data warehousing concepts that then leverage data mining techniques for reporting and analysis. These data feeds will then be pumped into systems that establish a baseline for performance and have built-in artificial intelligence that can detect anomalous behaviour indicative of a an instance of misuse or abuse, flagging and escalating the event accordingly.

B.Security Analysis

The purpose of an audit data reduction system is to reduce the overall cost and complexity associated with combining audit logs into one location and interface. These systems may have direct or indirect impact on the Confidentiality, Integrity or Availability of data or systems, depending on the source of the logs and the type of misuse or abuse detected. In general, ADR systems are a countermeasure designed to better detect instances of misuse or abuse. As the systems mature and further integrate with intrusion detection and analysis systems, the capability will also emerge to take protective and corrective actions. For example, intrusion detection and prevention systems (as will be discussed below) already have the capability to react dynamically and in real-time to detected threats. Using audit data reduction systems to accurately detect misuse or abuse in real-time holds the promise of integrating with these active response systems and thus extend its countermeasure capabilities.

V.FIREWALLS[10] [11] [12] [13] [14]

A firewall is defined as a "component or set of components that restricts access between a protected network and the Internet, or between other sets of networks." [15] Firewalls are network security resources that are defined to control the flow of data between two or more networks. From a high-level perspective, they can serve as a choke-point, designed to restrict, or choke, the flow of network traffic, or as a gateway that performs further processing on the traffic beyond simple choking restrictions. According to Zwicky, et al, firewalls can generally be placed into two categories: Packet Filters or Proxies. Per discussion in EMSE 218, these categories can be broadened to include circuit-level gateways and stateful inspection devices. Blanding[16] adds a third category of hybrid or complex gateways to Zwicky's initial pair.

In reality, the Blanding definition is probably the most correct in that firewalls either perform as a packet filter, a proxy, or as some combination of the two. Other types of firewall simply expand upon those original base types. For example, most proxies today have additional capabilities to perform content management at the application level, detecting inappropriate or unacceptable content, such as through a web or mail session.

[10] Manu, *Firewall Basics* (Unknown: SecurityDocs.com, 2004, accessed 06 November 2004); available from http://www.securitydocs.com/library/2413; Internet.
[11] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000).
[12] Simson Garfinkel and Gene Spafford, *Practical Unix & Internet Security, 2nd Edition* (Cambridge: O'Reilly, 1996).
[13] Lecture notes from EMSE 218, taken 20 October 2004.
[14] Purdue University, *Firewalls* (West Lafayette: CERIAS, undated, accessed 12 October 2004); available from http://www.cerias.purdue.edu/about/history/coast_resources/firewalls/; Internet.
[15] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p102.
[16] Steven F. Blanding, "Secured Connections to External Networks," in <I>Information Security Management Handbook, 4th Edition<I>, ed. Harold F. Tipton and Micki Krause (Boca Raton: Auerbach, 2000), p59-61.

Also, many firewalls provide capabilities like Network Address Translation (NAT) that provide a logical separation between networks by changing the underlying numbering scheme (IP addressing). NAT is an important feature because it allows organizations to interconnect their resources internally using IP address space that is reserved for internal use by RFC 1918. This reserved space is not routable on the Internet, and thus is not directly accessible to attackers outside the firewall performing the NAT.

A survey of various vendor web sites, such as Cisco, Checkpoint, NetScreen, CyberGuard, BlueCoat and Secure Computing, reflects the reality that most firewalls are now hybrids. This notion is further reinforced when reading through the Firewall Criteria v4.1[17] for ICSA Labs' Firewall Certification program. No firewall can receive a certification today without being aware of state, thus making it a stateful inspection firewall. However, basic firewalls, like those sold by Cisco, Checkpoint and NetScreen, are essentially just packet filtering, with the additional capabilities of tracking the state of a network session. Checkpoint extends this base design further by also providing some application-specific proxy components. CyberGuard, BlueCoat and Secure Computing, on the other hand, produce firewalls that are primarily proxies. Again, however, because of their adherence to the ICSA criteria, they also are aware of state, at least to some degree, and thus are able to perform basic packet filtering functions, too. Therefore, today, it is probably safe to say that there is only one kind of firewall, and that is a hybrid or complex gateway.

[17] http://www.icsalabs.com/html/communities/firewalls/certification/criteria/criteria_4.1.shtml

A.Business Analysis

The cost of a firewall today is minimal, and is greatly outweighed by the vast utility it serves. Firewalls need not be expensive solutions, but can be based on generic computer components that make use of free, open-source operating systems and software. Furthermore, these simple solutions do not require extensive and expensive hardware, but can oftentimes simply include a processor, memory and a storage device (like a CD-ROM). If the security requirements for an environment are stricter, then there are also many commercially viable solutions that range in price and capability. Several vendors sell firewalls of varying types that can handle a variety of network security needs. Whether those needs be for application proxies, or redundant packet filtering with automatic failover and recovery capabilities, or web proxies with content management capabilities to protect end-users against the hazards of unsafe web browsing, the only limitation today is in the size of the budget.

B.Security Analysis

"Firewalls are powerful tools, but they should never be used *instead* of other security measures. They should only be used *in addition* to such measures." [18] The primary role of a firewall, in the traditional sense, is to protect against unauthorized access of resources via the network as part of a “defense in depth” solution. This role serves to ensure the integrity of data and systems while also limiting the availability of those resources to malfeasants. Despite all the advances in firewall technology over the past 20 years, the fundamental role of the firewall has not changed. What has changed is the

[18] Simson Garfinkel and Gene Spafford, *Practical Unix & Internet Security, 2nd Edition* (Cambridge: O'Reilly, 1996), p637.

ability to integrate firewalls with other technologies, such as intrusion detection and analysis systems. Such integration can lead to providing an active response capability that blocks access to detected attackers in a real-time manner. Furthermore, in addition to serving in a protecting role, the audit and activity logs produced by a firewall can be used for detecting attacks, which can in turn result in the initiation of corrective actions, as has already been mentioned.

Firewalls, today, serve as a basic building block within security infrastructures. At the same time, as quoted above, they are not the "silver bullet" of information security. Implementation of a firewall is no guarantee of security and should be combined with the other security technologies described within this paper.

VI.INTRUSION DETECTION AND ANALYSIS SYSTEMS

The concept of intrusion detection has been around since 1980.[19] In its most essential form, intrusion detection is designed to detect misuse or abuse of network or system resources and report that occurrence. This detection occurs as a result of identifying behaviour based on anomalies or signatures. The most common form of intrusion detection system (IDS) today relies on signature-based detection.

The security industry has greatly expanded intrusion detection over the past years to incorporate several advanced concepts. Beyond basic detection and alerting, most systems today bill themselves as having "intrusion prevention" capabilities; otherwise

[19] Paul Innella, *The Evolution of Intrusion Detection Systems* (Unknown: SecurityFocus.com, 2001, accessed 12 October 2004); available from http://www.securityfocus.com/infocus/1514; Internet.

known as active response. The concept of intrusion prevention is that an activity can be detected reliably and then stopped, either at the host or network level, by the detecting system. From the network perspective, this response could be as simple as detecting an abusive TCP-based network connection and issuing a TCP Reset (RST) packet to both the source and destination hosts, forging the IP header information to impersonate each side.

Additionally, significant advances have been made in the areas of event correlation and anomaly detection. Event correlation is an approach wherein multiple alerts that may appear disparate are able to be linked together based on common criteria, such as time or method or target, and result in an escalated alert, if not a coordinated automatic response. Anomaly detection is similar to event correlation, though its primary role is to scientifically determine a baseline for performance, such as across a network or group of hosts, and then generate alerts when performance deviates significantly from that baseline.

The following sections discuss each of these technologies, providing an overview and then a respective business and security analysis.

A. Intrusion Detection Systems (IDS)[20]

Intrusion detection systems are typically classified according to their primary method of detection: network-based, host-based, hybrid, or network-node. Network-based detection

[20] Paul Innella, *The Evolution of Intrusion Detection Systems* (Unknown: SecurityFocus.com, 2001, accessed 12 October 2004); available from http://www.securityfocus.com/infocus/1514; Internet.

captures packets directly off the network, while host-based detection resides on a host and captures data as it flows into and out of that host. Hybrid systems aggregate the capabilities of network-based and host-based systems whereas network-node systems try to function like a network-based system while residing on a host.

Today, IDS has begun to mature to the point where most systems can be operated as a hybrid, if the business desires. The main approach used, such as through the open-source product Snort, is to conduct network- and/or host-based scanning using a signature set and then aggregate alerts to a single host for management of those alerts. More advanced systems have additional capabilities, as will be discussed in the following sections, such as intrusion prevention, anomaly detection, and event correlation.

Intrusion detection systems, as a whole, have a couple key limitations. First, they are typically limited in the same way that antivirus is limited in that successful detection is based on having a good signature that matches known bad traffic. With network detection, this signature limitation is particularly challenging because too literal of a string can result in a detection failure. Furthermore, IDS are limited by how much network traffic they can process in a given period of time. For example, most IDS today will claim to be able to monitor 1Gbps of traffic in real-time, though actual testing, such as in the IDS Lab at ICSA Labs, has proven that these products are actually often performing at much less than 1Gbps. Even worse, backbone network providers are often running at much higher speeds than 1Gbps, such as over OC-48 or OC-192 networks, which are 2.488 Gbps and 9.952 Gbps, respectively. This means that the needs and

expectations for performance and throughput are very high and not reasonably being met by commercial productions.

In addition to being limited by signatures and performance, most IDS also include management concerns with respect to the number of signatures being managed and the number of alerts being generated. Frustrations arising from these many limitations have led to advances in management of the base IDS, and will be discussed in the Anomaly Detection Systems and Event Correlation Systems sections below.

1.Business Analysis

Intrusion detection systems are still maturing as a product. Advances in event correlation, anomaly detection and active response have made their use much more appealing. However, the cost of deployment and management is still almost at a break-even point with the benefits derived. Networks that are particularly mature and clean have a much greater likelihood of reaping large benefits from an enhanced IDS deployment, whereas networks that are not well-designed and that are poorly managed will have a very difficult time tuning signatures to their environment and establishing performance baselines.

Quality IDS software is free through open-source initiatives such as Snort. Thanks to Snort, all a company really needs is a reasonably sized PC with one or more high-speed network cards and the know-how to install and manage the product on a compatible operating system, which may also be free. However, the open-source management tools that are available for use with Snort, such as ACID and SnortCenter, leave much to be desired and often force companies toward commercial solutions.

Most commercial solutions still tend to be rather expensive and require considerable training. One interesting development is the integration of intrusion detection solutions with firewall products, such as has been done by Cisco, Checkpoint and NetScreen. As will be discussed in the following section on intrusion prevention systems (IPS), this advance has allowed IDS to evolve to include active response capabilities, particularly from the network perspective.

Overall IDS has value for most organizations that have their network in good working order. However, understaffed and poorly architected environments will likely see IDS as an unacceptable hassle and cost. For those organizations, there are alternative solutions. Several security companies are in the market providing outsourced installation, maintenance and monitoring of IDS solutions. These "managed security solutions" providers may be beneficial for organizations that want the benefits of an IDS, even in a limited capacity, but that cannot afford to implement and manage the IDS themselves.

2.Security Analysis

The original role of IDS was to detect threats on networks and hosts. This role has evolved to include active response capabilities that allow it to protect resources and correct misuse or abuse on networks or hosts. IDS can today serve in a role that impacts Confidentiality, Integrity and Availability, depending on the signature set deployed, the effectiveness of alert management, and whether or not an active response capability exists.

B. Intrusion Prevention Systems (IPS)[21]

Intrusion prevention systems, or IPS, are often defined as "any device (hardware or software) that has the ability to detect attacks, both known and unknown, and prevent the attack from being successful." [22] IPS have grown from a desire to combine the deep-inspection capabilities of IDS with the blocking capabilities of firewalls. These blocking capabilities, often referred to as active response, allows the detection of a policy violation to be translated in real-time into a policy-based action designed to impede or stop the violation.

There are a few variations on IPS, but the most common is the inline network-based system. Another variation of IPS are the so-called "Layer 7 switches" that have matured to include DoS and DDoS detection and mitigation based on an awareness of traffic at the application layer of the OSI model. Also, host-based application firewalls have been integrated with IDS capabilities to allow for application-specific active response capabilities based on a general policy instead of a signature set. Hybrid switch solutions are network-based, but operate similar to the application firewalls.

All of these types of IPS have two things in common: they generate an alert, based either on a signature or a policy, and they initiate a response, as has been programmed into the system. These alerts may occur as the result of a signature match or a violation of a

[21] Neil Desai, *Intrusion Prevention Systems: the Next Step in the Evolution of IDS* (Unknown: SecurityFocus.com, 2003, accessed 12 October 2004); available from http://www.securityfocus.com/infocus/1670; Internet.
[22] Neil Desai, *Intrusion Prevention Systems: the Next Step in the Evolution of IDS* (Unknown: SecurityFocus.com, 2003, accessed 12 October 2004); available from http://www.securityfocus.com/infocus/1670; Internet.

security policy setup specific for an application, and the response may range from choking the flow of traffic to terminating or blocking the offending traffic altogether.

There are a couple key limitations to IPS, as exist for IDS. Those limitations include accurate detection, the ability to handle the full throughput of a network, and the ability to generate the response correctly and in a timely manner. The throughput issue has been discussed above. The matter of accuracy becomes increasingly important when discussing an active, automated response to a detected event. If proper and allowed traffic is incorrectly detected by a signature or as a policy violation, that traffic may be inappropriately subjected to the active response. In particular, known good traffic may be terminated or blocked, resulting in a negative impact to the business. As for generating the response correctly in a timely manner, this limitation pertains to the ability of the IPS to not only detect correctly, but to select the correct response based on a policy, and then be able to issue that response while the offense is still occurring. Choosing the proper response can become challenging when dealing with automated escalations.

1.Business Analysis

Most IDS systems today include some manner of IPS capabilities. Given a well-defined set of signatures or policies, it makes sense to deploy an IDS with IPS capabilities, particularly on the perimeter of your network, and in front of highly valuable assets. The cost of these systems is comparable to that discussed above in the IDS Business Analysis (VI.A.1). Ultimately, successful deployment and return on investment will relate directly to how well the network is architected, how well the solution is managed, and how much thought has gone into the overall security management of the organization.

2.Security Analysis

IPS expands the basic detection capabilities of IDS to include definite corrective capabilities. These corrective capabilities have the related benefit of protecting resources based on security policies. These capabilities work together to protect the Confidentiality, Integrity and Availability of systems and data.

*C.*Event Correlation Systems (ECS)[23]

Event Correlation Systems build on the successes of Intrusion Detection Systems by providing a better mechanism for aggregating, managing and correlating IDS events, such as are generated through signature detections or policy violations. ECS goes beyond simply pulling together event logs from IDS, however. ECS allows for the aggregation of log data from multiple sources, including firewalls, hosts, applications, and of course IDS. Most ECS solutions serve a dual role as a data warehouse for logs and by providing a data mining interface (manual and automated) to make use of the data stored in the warehouse.

The primary benefit of the Event Correlation System is in its ability to correlate events from multiple systems and generate smart alerts, along with the capability to escalate alerts, based on that correlation. Event Correlation Systems are usually comprised of

[23] Russell Kay, *Event Correlation* (Unknown: COMPUTERWORLD, 2003, accessed 12 October 2004); available from http://www.computerworld.com/networkingtopics/networking/management/story/0,10801,83396,00.html; Internet.

several key activities: Compression, Counting, Suppression, Generalization and Time-based correlation. These activities are best defined by Kay[24]:

> **Compression** takes multiple occurrences of the same event, examines them for duplicate information, removes redundancies and reports them as a single event. So 1,000 "route failed" alerts become a single alert that says "route failed 1,000 times."
>
> **Counting** reports a specified number of similar events as one. This differs from compression in that it doesn't just tally the same event and that there's a threshold to trigger a report.
>
> **Suppression** associates priorities with alarms and lets the system suppress an alarm for a lower-priority event if a higher-priority event has occurred.
>
> **Generalization** associates alarms with some higher-level events, which are what's reported. This can be useful for correlating events involving multiple ports on the same switch or router in the event that it fails. You don't need to see each specific failure if you can determine that the entire unit has problems.
>
> **Time-based correlation** can be helpful establishing causality -- for instance, tracing a connectivity problem to a failed piece of hardware. Often more

[24] Russell Kay, *Event Correlation* (Unknown: COMPUTERWORLD, 2003, accessed 12 October 2004); available from http://www.computerworld.com/networkingtopics/networking/management/story/0,10801,83396,00.html; Internet.

information can be gleaned by correlating events that have specific time-based relationships. Some problems can be determined only through such temporal correlation.

1.Business Analysis

ECS is the solution that is most desirable and has the potential for the biggest return on investment. However, implementation of such a system has proven to be very challenging for vendors. As a result, these systems tend to be very expensive and not terribly reliable. Instead, the Anomaly Detection approach, as discussed below, has been conceived and is beginning to receive increased market share. In the future, it is hoped that ECS will mature to the point where it can be integrated to round-out the Intrusion Detection and Analysis System.

2.Security Analysis

The primary function of ECS is to better detect events within the enterprise. Once reliable detection occurs, then other capabilities, such as active response, can be developed with it. Until that time, however, this solution is primarily aimed at protecting the Integrity of systems and data as a result of detecting active threats against them.

*D.*Anomaly Detection Systems (ADS)[25] [26]

[25] Christina Yip Chung, *Anomaly Detection in Database Systems* (Davis: UC Davis Computer Security Laboratory, 1999, accessed 12 October 2004); available from http://seclab.cs.ucdavis.edu/projects/anomaly.html; Internet.
[26] Roy A. Maxion and Kymie M.C. Tan, *Benchmarking Anomaly-Based Detection Systems* (Pittsburgh: Carnegie Mellon University, 2000, accessed 12 October 2004); available from http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/maxion/www/pubs/maxiontan00.pdf; Internet.

Anomaly Detection Systems are an extension of Intrusion Detection Systems (or Misuse Detection Systems, as defined by Chung). Per Maxion and Kymie, "[anomaly] detection is a key element of intrusion detection and other detection systems in which perturbations of normal behavior suggest the presence of intentionally or unintentionally induced attacks, faults, defects, etc." [27] This type of detection is based largely on the rules of probability and predictable, taking into consideration log data from multiple sources (much as is done in ECS), but applying theories of predictability to these logs and automatically generating a best guess as to whether or not a misuse, or abuse, is occurring. In its basest form, ADS generates a baseline for performance and then monitors for behaviour that deviates from that baseline. In its more advanced, optimized form, ADS dynamically calculates the current performance based on aggregate log data and determines whether or not the current level of performance is deviant from expected levels.

As outlined in Maxion and Kymie, one of the key challenges to ADS is in performance. A large number of calculations must be performed on the fly to determine whether or not the aggregate logs can be correlated and weighted in such a manner as to predict an instance of misuse. Maxion and Kymie theorized that the type and source of data used by an ADS could have an impact on its performance. This theory was proven through their experimentation, indicating, then, that ADS is subject to performance variation, depending on data and source factors. Unfortunately, performance variance is not

[27] Roy A. Maxion and Kymie M.C. Tan, *Benchmarking Anomaly-Based Detection Systems* (Pittsburgh: Carnegie Mellon University, 2000, accessed 12 October 2004); available from http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/maxion/www/pubs/maxiontan00.pdf; Internet.

something generally appreciated by businesses and could stand to limit its adoption within corporate security environments.

1.Business Analysis

Anomaly detection systems are an emerging solution related in part to intrusion (or misuse) detection systems and event correlation systems. This reality as an emerging technology limits the number of commercial solutions available and increases the cost of deployment. Some organizations have gone so far as to develop rudimentary ADS solutions in-house in order to defer commercial expenses. However, the overall value of these systems is limited by the primitive routines performed.

Ultimately, ADS and ECS represent the ideal solutions that will maximize return on investment for detection of threats within a security infrastructure. Once solutions begin to mature, competition emerges, and prices begin to drop, we will likely see a wide adoption of these types of solutions. Until that time, only the largest organizations, with the necessary resources to implement such a solution, will likely see the utility of ADS or ECS. Small and medium sized organizations will likely need to be content with basic IDS and IPS capabilities for the foreseeable future, banning a major breakthrough in performance and reliability that can reduce the overall total cost of ownership while maximizing the value. Integration of these solutions with active response capabilities and firewalls will continue to mature as the core products themselves mature.

2.Security Analysis

ADS are primarily designed to detect threats to the organization. This detect capability may be expanded in the future to include protect and correct capabilities, but only after the product has matured further. The general goal of ADS, as is true with most intrusion detection related solutions, is to primarily ensure Integrity, with secondary goals of ensuring Availability and Confidentiality. Detection can be used universally to ensure all three aspects of the CIA approach.

VII.NETWORK MAPPING

Network mapping is defined as "the study of the physical connectivity of the Internet."[28] In its most common form, network mapping is used to document the layout of a local area network (LAN) as part of an overall security assessment. This use is a form of intelligence gathering and oftentimes precedes the actual assessment of targeted systems.

Network mapping has evolved over the years from the simple performance of "PING" or "CONNECT" attempts to more extensive and subversive (or "quiet") methods of detection. Today, the most popular tool for performing network mapping is the open-source tool Nmap.[29] Nmap is capable of testing for the presence of nodes on a network based on a variety of detection techniques, including the use of Internet Protocol (IP), Transmission Control Protocol (TCP) and Universal Datagram Protocol (UDP). Each of these protocols has a unique flavor, and thus can generate varying results. Furthermore,

[28] Wikipedia, *Network Mapping* (St. Petersburg: Wikipedia, 2004, accessed 12 October 2004); available from http://en.wikipedia.org/wiki/Network_Mapping; Internet.
[29] Fyodor, *Nmap Security Scanner* (Unknown: Insecure.org, undated, accessed 12 October 2004); available from http://www.insecure.org/nmap/index.html; Internet.

Nmap has additional capabilities for subverting network security devices like firewalls and intrusion detection systems. It can take as input a host name, an IP address, a range of IP addresses, or a network or subnetwork. It may also take configurable parameters of "dummy" source addresses to help camouflage to network sensors what it is trying to do.

The goal of network mapping is to determine would nodes are active on a network. This basic determination can be developed further to identify how far away the nodes are from the scanning host. Operating system identification may also be performed by tools like Nmap, though this functionality is an extension of network mapping and not core to its capabilities.

A.Business Analysis

Network mapping is a cheap and valuable tool for reviewing the existence of nodes on a network. Running a network mapping tool on a regular basis and comparing its results can assist an organization in ensuring that no nodes are being added to the network without proper authorization. Since the most popular tool, Nmap, is free and has been ported to many operating systems, including Linux, UNIX, Windows and Mac OS, the only real costs are in terms of performance and processing.

There are a couple potential risks and limitations for network mapping. First, some applications and systems do not respond well to probes from network mapping tools. Mainframes, for example, have been known to respond poorly to raw network socket requests. Thus, network mapping could cause instability in a mainframe, or at least

generate a large number of alerts. Additionally, network mapping can be limited by certain types of network and firewall rules. Whereas network mapping used to be able to circumvent firewalls using various packet manipulation techniques, most firewalls today are aware of state and thus effectively block circumvention. Additionally, intrusion detection systems, which may also be circumvented, have the capability today to be tuned so as to more optimally detect the occurrence of network mapping.

B.Security Analysis

Network mapping is a form of detection, from the standpoint that it detects nodes on a network, which can in turn be used to determine whether or not a given node is authorized to be on the network. Network mapping may also be construed as a form of protection, since the actions that derive from comparing network mapping data sets could result in removal of unauthorized nodes from the network.

From the standpoint of Confidentiality, Integrity and Availability, network mapping primarily serves the goal of ensuring the Integrity of the network. It may also be used to verify that certain nodes remain available on a network. Network mapping does not have any impact on Confidentiality, unless one were to spin the impact along the following line: a node, such as an IDS sensor, is placed on the network and configured so as not to be detectable by network mapping; however, a misconfiguration results in causing the sensor to respond to network mapping requests, revealing its location, and possibly its identity; thus, network mapping can ensure the confidentiality of “hidden” network nodes.

VIII.PASSWORD CRACKING

According to Wikipedia, "[password] cracking is the process of recovering secret passwords stored in a computer system."[30] Password cracking may serve to recover a lost password or to compromise an unknown password for the purposes of gaining unauthorized access to a system or data. Additionally, password cracking may be used as a preventative measure to ensure that strong passwords are being used by system users.

Most passwords today are maintained as a hashed, rather than encrypted, value. Hashing means taking a password string and using it as an input for an algorithm that results in an output that does not resemble the original input. Unlike encryption, hashing only works one way and cannot be decrypted. Hashing passwords before storing them is far more efficient than encrypting and decrypting passwords on the fly. Thus, when a user attempts to login, their submitted password is hashed, and the hashed value is compared with the hashed value stored on the system. Given an exact hash match, the login is approved and the user is considered authenticated.

The best commercial use of password cracking is as a preventative measure, ensuring that users are choosing high quality (or strong) passwords. According to @stake, maker of the popular l0phtcrack password cracking utility, "experts from SANS, industry, government, and academia cite weak passwords as one of the most critical security threats

[30] Wikipedia, *Password cracking* (St. Petersburg: Wikipedia, 2004, accessed 12 October 2004); available from http://en.wikipedia.org/wiki/Password_cracking; Internet.

to networks."[31] In the current context, passwords are the primary method for authentication, despite the availability of better solutions, as described in Section II above. Thus, protection of passwords and ensuring strong passwords against simple attacks is of the utmost importance.

Passwords are typically subjected to a combination of two kinds of attacks: brute-force and dictionary (or word-list). Brute-force attacks attempt to iterate through every possible password option available, either directly attempting to the test password against the system, or in the case of a captured password file, comparing the hashed or encrypted test password against the hashed or encrypted value in the file. In a dictionary attack, a list of common passwords, oftentimes consisting of regular words, is quickly run through and applied in a similar manner as with the brute-force attack.

Dictionary attacks are oftentimes very effective unless systems require users to choose strong passwords. For example, the maintainers of the popular open-source password cracking tool John the Ripper sell collections of word lists on CD. The CDs include word lists for more than 20 human languages, plus common and default passwords and unique words for all combined languages. For around $50 an individual wanting to execute a massive dictionary-based attack could have access to over 600MB of word list data.[32] The ready availability of such data sets for use in dictionary attacks means that, unless a strong password is selected, it is very likely that the password can be cracked in a

[31] @stake, *@stake LC 5* (Cambridge: @stake, undated, accessed 12 October 2004); available from http://www.atstake.com/products/lc/; Internet.
[32] Openwall Project, *John the Ripper password cracker* (Moscow: Openwall, undated, accessed 12 October 2004); available from http://www.openwall.com/john/; Internet.

reasonable amount of time. This is especially true of passwords that are based on human-readable words.

A strong password is most often defined as a string of eight (8) or more characters that mix upper- and lower-case letters, numbers and special characters. Strong passwords do not resemble words, and are best when generated at random.[33] One suggested approach is picking a passphrase and either using the passphrase in its entirety or picking the leading letters from each word in the phrase and substituting numbers and special characters for some of the letters. Certain password hashing algorithms produce stronger hash values with longer passwords while others produce stronger hash values based on increased complexity of the password.

In addition to requiring users to choose strong passwords, it is also incumbent upon system administrators to require that passwords be changed frequently. Conventional wisdom indicates that no password should have a lifetime greater than 90 days, and for highly critical systems the lifetime should be 30 days or less. One exception to this rule involves two-factor authentication where a password is coupled with a stronger authentication method, such as tokens or biometrics.

A.Business Analysis

Passwords hold a prevalent place within the security infrastructure throughout most, if not all, organizations. Until passwords are replaced by stronger forms of authentication,

[33] A. Cliff, *Password Crackers - Ensuring the Security of Your Password* (Unknown: SecurityFocus.com, 2001, accessed 12 October 2004); available from http://www.securityfocus.com/infocus/1192; Internet.

such as tokens or biometrics, it is absolutely necessary that the use of strong passwords be enforced. Therefore, the benefit of buying word lists and password cracking software and running them regularly, particularly on key systems, greatly outweighs the costs. One downside is where centralized authentication has not been implemented. In those cases, while it is likely that users will use the same password across multiple systems, the cost in time of running password cracking against all systems becomes challenging. Thus, in addition to password cracking, it is also useful to implement a centralized authentication system that results in fewer password files to test.

B.Security Analysis

Password cracking is primarily a protective countermeasure. It is designed to ensure that passwords used in various authentication mechanisms are strong enough to prevent casual dictionary-based attacks. It is assumed, however, that a brute-force attack can be 100% successful given enough time. As such, it is vitally import to combine password cracking with strict systematic requirements for strong passwords and regular password rotation. Password cracking helps ensure the Confidentiality and Integrity of data and systems by propping-up the authentication system.

IX.PUBLIC KEY INFRASTRUCTURE[34]

[34] The following general resources are available, but not quoted in this paper :

- several PKI links: http://www.pki-page.org/
- more PKI docs: http://www.opengroup.org/public/tech/security/pki/
- x509 WG http://www.ietf.org/html.charters/pkix-charter.html
- Federal PKI Steering Committee: http://www.cio.gov/fpkisc/
- PKI and the Law: http://www.pkilaw.com/

Public Key Infrastructure was once thought to be the silver bullet for solving security and privacy on the Internet, as well as providing a framework for secure business transactions across shared network resources. The reality is that PKI is complex, expensive, and very difficult to implement well. Clarke has gone so far as to claim, with significant proof, that PKI will remain a failure and offers alternatives that seek to improve or supplant the current X.509 standard for PKI. "Its key deficiencies are its inherently hierarchical and authoritarian nature, its unreasonable presumptions about the security of private keys, a range of other technical and implementation defects, confusions about what it is that a certificate actually provides assurance about, and its inherent privacy-invasiveness."[35]

According to Wikipedia, a Public Key Infrastructure is "an arrangement, usually carried out by software at a central location together with other coordinated software at distributed locations, which provides for third party (often termed a trusted third party) vetting of and vouching for user identities and for binding of public keys to users (typically in certificates) and vice versa."[36] The most common form of PKI today is Secure Socket Layer (SSL) certificates used through the Internet for securing web browsing sessions. Companies such as VeriSign and Thawte make available servers on the Internet through which another organization's SSL certificate can be verified by a client web browser as being authentic and non-revoked.

[35] Roger Clarke, *Conventional Public Key Infrastructure: An Artefact Ill-Fitted to the Needs of the Information Society* (Canberra : Clarke, 2000, accessed 12 October 2004); available from http://www.anu.edu.au/people/Roger.Clarke/II/PKIMisFit.html; Internet.
[36] Wikipedia, *Public key infrastructure* (St. Petersburg: Wikipedia, 2004, accessed 06 November 2004); available from http://en.wikipedia.org/wiki/Public_key_infrastructure; Internet.

In more complex scenarios, PKI can be deployed internally to an organization for various purposes, such as secure internal communication, providing encryption services to data and systems, digitally signing code, and providing encryption materials allowing users to digitally sign communication. Typically, though, enterprise PKI solutions are provided primarily as part of an authentication system to better prove and secure an individual's identity.

Wikipedia provides a decent overview of PKI and its history. [37] Additionally, the National Institute of Standards and Technology (NIST) has stepped-up to provide public leadership for deployment and support of PKI in federal environments, as well as to help steer the development and standardization of associated technologies.[38]

A.Business Analysis

PKI has historically been considered a pipedream that will never come to fruition. Considerable criticism has been levied against it due to the associated cost and complexity. The major deployments of PKI today seem to focus around supporting SSL for Internet transactions. However, PKI has finally begun to evolve and mature, to the point where other large organizations have decided to install enterprise solutions. For example, AOL now has its own public PKI that it can use for the purposes of generating SSL certificates, among other things. Proof of this deployment can be found by reviewing the root certificates issued with all major browsers.

[37] Wikipedia, *Public key infrastructure* (St. Petersburg: Wikipedia, 2004, accessed 06 November 2004); available from http://en.wikipedia.org/wiki/Public_key_infrastructure; Internet.
[38] National Institute of Standards and Technology, *NIST PKI Program* (Washington: NIST, 2004, accessed 12 October 2004); available from http://csrc.nist.gov/pki/; Internet.

The use of PKI for enhancing authentication and identification is a laudable goal, but one that is very expensive in achieving. This goal could also be achieved through use of cheaper authentication systems combined with other authentication methods, such as tokens or biometrics. Thus, the utility of a PKI for the average organization appears to be minimal, and the cost is generally very prohibitive.

The types of companies that would benefit from deploying their own PKI include large Internet companies that not only conduct a lot of business across the Internet, but that provide business services to other organizations across the Internet. Additionally, large software development companies may benefit from implementing a PKI in order to sign code and applications to help assure the integrity of their products and intellectual property. Federal agencies and financial services companies may also benefit from deploying some form of PKI in order to establish a network of trust upon which customers and citizens can rely.

In the end, however, the extreme cost of successfully implementing a PKI solution generally outweighs most of the benefits that may be derived. In the case of small or medium sized software development companies, it may in fact be cheaper to rely on code signing from a trusted third party rather than to conduct the code signing with an in-house PKI.

B.Security Analysis

The main role of PKI as a countermeasure is to protect against attack and compromise. Whether it be integrated into an authentication system or part of a code signing system, the overall goal is to ensure Integrity. Additionally, PKI can serve in a capacity of ensuring that Confidentiality of data through trusted encryption mechanisms that leverage trusted encryption materials. PKI may, however, have a negative affect on the Availability of data or systems. If the PKI fails, the associated materials or mechanisms may not function properly to decrypt data, or to allow for proper authentication to occur. Since a secure system will "fail safe," failure of the PKI should fail to a closed state that disallows access, but in turn impacting Availability.

X. VIRTUAL PRIVATE NETWORKS[39]

A Virtual Private Network (VPN) is a private communications network that makes use of public networks, oftentimes for communication between different organizations.[40] A VPN is not inherently secure, though in its most common incarnation it does utilize encryption to ensure the confidentiality of data transmitted. The VPN is often seen as a cheaper solution for deploying a private network than private leased-lines.[41] [42] They often serve to protect and ensure the integrity of communications[43] and may also protect the confidentiality of those communications when utilizing encryption.

Aside from the cost factor, VPNs have two main advantages: they may provide overall encryption for communications and they allow the use of protocols that are otherwise difficult to secure. [44] In contrast, Zwickey sites the two main disadvantages of VPNs being the reliance on "dangerous" public networks and extending the network that is being protected.[45]

There are three types of VPNs available today: dedicated, SSL and opportunistic. Dedicated VPNs, either in a gateway-to-gateway or client-to-gateway configuration,

[39] About.com has several links on VPNs that may be worth reviewing. http://compnetworking.about.com/od/vpn/
[40] Wikipedia, *Virtual private network* (St. Petersburg: Wikipedia, 2004, accessed 06 November 2004); available from http://en.wikipedia.org/wiki/Virtual_private_network; Internet.
[41] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p104.
[42] Robert Moskowitz, *What Is A Virtual Private Network?* (Unknown: CMP, undated, accessed 12 October 2004); available from http://www.networkcomputing.com/905/905colmoskowitz.html; Internet.
[43] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p119.
[44] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p120.
[45] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p121.

appear to currently be the most prominent deployment. However, SSL VPNs are increasing in popularity, serving as a lightweight, platform-independent client-to-gateway protection mechanism. Additionally, the concept of opportunistic encryption, as used with VPNs, was first posited in 2001 by the FreeS/WAN project, who's mission was to provide free standards-based VPN software under an open-source initiative. The concept of opportunistic encryption (OE) hinged on the notion that a VPN did not need to be in an "up" state at all times, but rather only needed to be activated when communication was occurring. Thus, gateways across the Internet could be configured to support encryption on an as-needed basis and would only have to setup the VPN when a connection from/through an OE-aware gateway was initiated. This model is similar to the traditional use of SSL on the Internet, except that instead of simply encrypting the traffic at the application layer, the encryption was actually occurring at the network and/or transport layer, and all happening transparent to the end-user.[46] The goal of implementing opportunistic encryption within free IPSEC-based VPNs was to transparently encrypt all Internet traffic.

Most virtual private networks today make use of IPSEC encryption. IPSEC provides network-level security for the Internet Protocol (IP) and is an extension of the original IPv4 standard. IPSEC makes use of the management and security protocol ISAKMP/Oakley and has the benefit of protecting against man-in-the-middle attacks

[46] Henry Spencer and D. Hugh Redelmeier, *Opportunistic Encryption* (Unknown: Freeswan.org, 2001, access 07 November 2001); available from http://www.freeswan.org/freeswan_trees/freeswan-1.91/doc/opportunism.spec; Internet.

during connection setup. IPSEC includes a number of other features, such as being usable by tunneling protocols.[47]

A.Business Analysis

Virtual private networks have a legitimate use in the business environment, especially when used in a secure manner, leveraging available encryption options. Given the growing prevalence and availability of cheap Internet access, a VPN can be used to securely and reliably replace more expensive leased lines. This replacement is particularly nice in environments where the data being transmitted is sensitive, but where interruption of connectivity will not represent a major disruption to the business.

Many hardware and software solutions are available today, with costs ranging from free (FreeS/WAN) to expensive (dedicated hardware-based solutions targeting high throughput). Most inexpensive networking equipment, such as the Linksys and Netgear lines of home user security devices, now support IPSEC-based VPNs.

B.Security Analysis

The basic goal of a Virtual Private Network is to ensure the integrity of the connection and communications.[48] When encryption is added, the goal of preserving confidentiality may also be achieved. One downside to VPNs is that they tend to be built on complex

[47] Robert Moskowitz, *What Is A Virtual Private Network?* (Unknown: CMP, undated, accessed 12 October 2004); available from http://www.networkcomputing.com/905/905colmoskowitz.html; Internet.
[48] Elizabeth D. Zwicky and others, *Building Internet Firewalls, 2nd Edition* (Cambridge: O'Reilly, 2000), p119.

systems and are prone to easy disruption, reducing the overall availability of data and communications.

From the perspective of countermeasures, the VPN primarily serves to protect data, though it may also dynamically correct. If logging is enabled and monitored, then attacks against the VPN may also result in meeting the need of detection, though that would be ancillary.

XI.VULNERABILITY SCANNING SYSTEMS

Vulnerability scanning is the "automated process of proactively identifying vulnerabilities of computing systems in a network in order to determine if and where a system can be exploited and/or threatened."[49] Vulnerability scanning typically relies on a handful of tools that identify hosts and then proceed to test them for known weaknesses. The automated scanning process should include three high-level steps: receiving authority to scan, determining the scope of the program, and establishing a security baseline (based on the number of vulnerabilities found per number of hosts scanned).[50] Additionally, a good vulnerability scanning program will securely manage the results of the scans and will have a proven plan and process in place for remediation of vulnerabilities that are uncovered. Vulnerability scanning should occur as part of an overall risk management framework, not as a standalone security countermeasure.

---

[49] Webopedia, *vulnerability scanning* (Darien: Jupitermedia, undated, accessed 12 October 2004); available from http://www.webopedia.com/TERM/V/vulnerability_scanning.html; Internet.

[50] Christopher Cook, *Managing Network Vulnerabilities in a DOE/NNSA Environment* (Kansas City: DOE, undated, accessed 12 October 2004); available from http://cio.doe.gov/Conferences/Security/Presentations/CookC.pps; Internet.

The most popular vulnerability scanning tool available today is also free, open-source software. Nessus[51] has become the de facto tool for vulnerability scanning over the past five (5) years, replacing commercial tools like CyberCop Scanner (discontinued), ISS Security Scanner, and eEye Retina. Vulnerability scanning has been around since the late 80s or early 90s, pioneered by Dan Farmer, co-author of the COPS[52] security tool. Originally, vulnerability scanning was host-based in nature, as COPS and TIGER were, but eventually expanded to include network-based scanning. There are still host-based scanners available, such as the Center for Internet Security's benchmark security tool[53]. More often, though, vulnerability scanning today is network-based.

Chapple provides a nice overview of the Nessus scanner and why it's preferable to its competition:

> "The Nessus tool works a little differently than other scanners. Rather than purporting to offer a single, all-encompassing vulnerability database that gets updated regularly, Nessus supports the Nessus Attack Scripting Language (NASL), which allows security professionals to use a simple language to describe individual attacks. Nessus administrators then simply include the NASL descriptions of all desired vulnerabilities to develop their own customized scans."[54]

[51] http://www.nessus.org/
[52] http://www.fish.com/cops/overview.html
[53] http://www.cisecurity.com/
[54] Mike Chapple, *Vulnerability scanning with Nessus* (Unknown: TechTarget.com, 2003, accessed 12 October 2004); available from http://searchsecurity.techtarget.com/tip/0,289483,sid14_gci938271,00.html?track=NL-20; Internet.

A.Business Analysis

As was the case with password cracking in Section VIII above, vulnerability scanning is a very cheap and useful practice. When conducted regularly and carefully, the use of an automated vulnerability scanning tool can provide considerable information about the overall risk landscape of technologies throughout an enterprise. Vulnerability scanning is particularly important for ensuring that Internet-accessible resources are properly secured before deployment, and to ensure that they remain secure after deployment.

Because the most common tools for conducting vulnerability scans is free, open-source software, there is very little reason not to make use of it. Furthermore, the installation and operation of a tool like Nessus does not require much technical acumen. More importantly, the information that can be gathered from the assessment can be invaluable. Operation of a basic vulnerability scanner is not complex. Making matters even better, tools like Nessus are thoroughly documented on the Internet and can often be found in pre-packaged bootable environments.

B.Security Analysis

Vulnerability scanning can contribute to countermeasures in all three areas of protect, detect and correct. The primary role of the scanning is to detect vulnerabilities in systems, but when used properly it will also contribute to protecting resources from being deployed insecurely and by providing adequate information to allow system administrators to correct vulnerabilities.

From the standpoint of Confidentiality, Integrity and Availability, vulnerability scanning most affects the Integrity of systems, though there may be ancillary benefits to Confidentiality and Availability. In detecting and resolving weaknesses in a system, the integrity of the system can be assured. Furthermore, ensuring the integrity of a system will help prevent the system from becoming compromised, resulting in a loss of confidentiality, or from being overly susceptible to attacks that may result in denying the availability of the system or associated application.

REFERENCES

1. @stake. *@stake LC 5*. Cambridge: @stake, undated, accessed 12 October 2004; available from http://www.atstake.com/products/lc/; Internet.

2. Blanding, Steven F. "Secured Connections to External Networks," in *Information Security Management Handbook, 4th Edition*, ed. Harold F. Tipton and Micki Krause. Boca Raton: Auerbach, 2000.

3. Chapple, Mike. *Vulnerability scanning with Nessus*. Unknown: TechTarget.com, 2003, accessed 12 October 2004; available from http://searchsecurity.techtarget.com/tip/0,289483,sid14_gci938271,00.html?track=NL-20; Internet.

4. Clarke, Roger. *Conventional Public Key Infrastructure: An Artefact Ill-Fitted to the Needs of the Information Society*. Canberra : Clarke, 2000, accessed 12 October 2004; available from http://www.anu.edu.au/people/Roger.Clarke/II/PKIMisFit.html; Internet.

5. Cliff, A. *Password Crackers - Ensuring the Security of Your Password*. Unknown: SecurityFocus.com, 2001, accessed 12 October 2004; available from http://www.securityfocus.com/infocus/1192; Internet.

6. Cook, Christopher. *Managing Network Vulnerabilities in a DOE/NNSA Environment*. Kansas City: DOE, undated, accessed 12 October 2004; available from http://cio.doe.gov/Conferences/Security/Presentations/CookC.pps; Internet.

7. Desai, Neil. *Intrusion Prevention Systems: the Next Step in the Evolution of IDS*. Unknown: SecurityFocus.com, 2003, accessed 12 October 2004; available from http://www.securityfocus.com/infocus/1670; Internet.

8. eBCVG IT Security. *Heuristic Scanning - Where to Next?*. Tel-Aviv: eBCVG, 2004, accessed 12 October 2004; available from http://www.ebcvg.com/articles.php?id=264; Internet.

9. Fyodor. *Nmap Security Scanner*. Unknown: Insecure.org, undated, accessed 12 October 2004; available from http://www.insecure.org/nmap/index.html; Internet.

10. Garfinkel, Simson and Gene Spafford, *Practical UNIX & Internet Security, 2nd Edition*. Cambridge: O'Reilly, 1996.

11. Innella, Paul. *The Evolution of Intrusion Detection Systems*. Unknown: SecurityFocus.com, 2001, accessed 12 October 2004; available from http://www.securityfocus.com/infocus/1514; Internet.

12. Kanish, Bob. *An Overview of Computer Viruses and Antivirus Software*. Unknown: Kanish, 1996, accessed 12 October 2004; available from http://www.hicom.net/~oedipus/virus32.html; Internet.

13. Kay, Russell. *Event Correlation*. Unknown: COMPUTERWORLD, 2003, accessed 12 October 2004; available from http://www.computerworld.com/networkingtopics/networking/management/story/0,10801,83396,00.html; Internet.

14. Manu. *Firewall Basics*. Unknown: SecurityDocs.com, 2004, accessed 06 November 2004; available from http://www.securitydocs.com/library/2413; Internet.

15. Maxion, Roy A. and Kymie M.C. Tan. *Benchmarking Anomaly-Based Detection Systems*. Pittsburgh: Carnegie Mellon University, 2000, accessed 12 October 2004; available from http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/maxion/www/pubs/maxiontan00.pdf; Internet.

16. Moskowitz, Robert. *What Is A Virtual Private Network?*. Unknown: CMP, undated, accessed 12 October 2004; available from http://www.networkcomputing.com/905/905colmoskowitz.html; Internet.

17. National Institute of Standards and Technology. *NIST PKI Program*. Washington: NIST, 2004, accessed 12 October 2004; available from http://csrc.nist.gov/pki/; Internet.

18. National Institute of Standards and Technology. *NIST Planning Report 02-1: Economic Impact Assessment of NIST's Role-Based Access Control (RBAC) Program*. Washington: NIST, 2002, accessed 12 October 2004; available from http://csrc.nist.gov/rbac/rbac-impact-summary.doc; Internet.

19. Openwall Project. *John the Ripper password cracker*. Moscow: Openwall, undated, accessed 12 October 2004; available from http://www.openwall.com/john/; Internet.

20. Purdue University. *CERIAS: Audit Trail Reduction Group*. West Lafayette: CERIAS, undated, accessed 12 October 2004; available from http://www.cerias.purdue.edu/about/history/coast/projects/audit-trails-reduce.php?output=printable; Internet.

21. Purdue University. *Firewalls*. West Lafayette: CERIAS, undated, accessed 12 October 2004; available from http://www.cerias.purdue.edu/about/history/coast_resources/firewalls/; Internet.

22. Richards, Donald R. "Biometric Identification," in <I>Information Security Management Handbook, 4th Edition<I>, ed. Harold F. Tipton and Micki Krause. Boca Raton: Auerbach, 2000.

23. Rotchke, Ben. *Access Control Systems & Methodology*. New York: SecurityDocs.com, 2004, accessed 06 November 2004; available from http://www.securitydocs.com/go/69; Internet.

24. Spencer, Henry and D. Hugh Redelmeier, *Opportunistic Encryption*. Unknown: Freeswan.org, 2001, access 07 November 2001; available from http://www.freeswan.org/freeswan_trees/freeswan-1.91/doc/opportunism.spec; Internet.

25. Tipton, Harold F. and Micki Krause. *Information Security Management Handbook, 4th Edition*. Boca Raton: Auerbach, 2000.

26. Webopedia. *vulnerability scanning*. Darien: Jupitermedia, undated, accessed 12 October 2004; available from http://www.webopedia.com/TERM/V/vulnerability_scanning.html; Internet.

27. Wikipedia. *Anti-virus software*. St. Petersburg: Wikipedia, 2004, accessed 06 November 2004; available from http://en.wikipedia.org/wiki/Anti-viral_software; Internet.

28. Wikipedia. *Computer virus*. St. Petersburg: Wikipedia, 2004, accessed 06 November 2004; available from http://en.wikipedia.org/wiki/Computer_virus; Internet.

29. Wikipedia. *Network Mapping*. St. Petersburg: Wikipedia, 2004, accessed 12 October 2004; available from http://en.wikipedia.org/wiki/Network_Mapping; Internet.

30. Wikipedia. *Password cracking*. St. Petersburg: Wikipedia, 2004, accessed 12 October 2004; available from http://en.wikipedia.org/wiki/Password_cracking; Internet.

31. Wikipedia. *Public key infrastructure*. St. Petersburg: Wikipedia, 2004, accessed 06 November 2004; available from http://en.wikipedia.org/wiki/Public_key_infrastructure; Internet.

32. Wikipedia. *Virtual private network*. St. Petersburg: Wikipedia, 2004, accessed 06 November 2004; available from http://en.wikipedia.org/wiki/Virtual_private_network; Internet.

33. Yip Chung, Christina. *Anomaly Detection in Database Systems*. Davis: UC Davis Computer Security Laboratory, 1999, accessed 12 October 2004; available from http://seclab.cs.ucdavis.edu/projects/anomaly.html; Internet.

34. Zwicky, Elizabeth D., S. Cooper and D. B. Chapman. *Building Internet Firewalls, 2nd Edition*. Cambridge: O'Reilly, 2000.