GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at http://www.giac.org/registration/gsec

Title: An Overview of Defense in Depth at each layer of the TCP/IP Model

Name: Jose A. Dominguez
Admin Ver: 2.3

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

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

1. INTRODUCTION **3**

2. WHAT IS TCP/IP? **3**

3. WHAT IS DEFENSE IN DEPTH? **4**

4. LINK LAYER **5**

4.1 PHYSICAL SECURITY 6
4.2 MEDIA ACCESS CONTROL (MAC) LAYER FILTERING 6
5. NETWORK LAYER 7
5.1 FIREWALLS 8
5.2 ACCESS CONTROL LISTS 8
5.3 VIRTUAL PRIVATE NETWORK 9

6. TRANSPORT LAYER **9**

6.1 TRANSMISSION CONTROL PROTOCOL (TCP) 9
6.2 USER DATAGRAM PROTOCOL (UDP) 10
6.3 SECURE SOCKET LAYERS (SSL) 10

7. APPLICATION LAYER **10**

7.1 APPLICATION PROXY FIREWALLS 11
7.2 HOST BASED FIREWALLS 11
7.3 ANTI VIRUS SCANNER 12

8. CONCLUSION **12**

9. REFERENCES **14**

9.1 BOOKS/GUIDES 14
9.2 PAPERS/LINKS 14

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

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

An Overview of Defense in Depth at each layer of the TCP/IP Model

1. Introduction

The goal of this paper is to discuss general security concepts with a focus on security at each layer of the TCP/IP (Transmissions Control Protocol/Internet Protocol) model while providing a fundamental and conceptual understanding of defense in depth. It is by no means an all-inclusive reference to securing an environment.

Enabling access to critical applications and data while maintaining confidentiality, integrity and availability can be a difficult task. Although an isolated computer is useful, it becomes much more powerful when connected to a web of computers (i.e., a network). The greatest advantage of networking a computer is sharing information. Why maintain copies of applications and data on multiple systems when one system can house them and others can view and use them? The most well known network is the Internet, which is actually made up of hundreds of thousands of smaller networks with the ability to communicate with one another through routers. As with human language, specific rules must be followed in order for communication between the networks to take place. These rules are called protocols and in the case of the Internet, this common language is known as the TCP/IP Protocol suite.

As mentioned, the advantage of networking a system is to share information; however, you should always keep in mind the risks associated with creating networks. To ensure confidentiality, integrity, and availability of data, you will want to limit access to certain information and the systems that house that information. Therefore, when contemplating network security, it is important to begin your thought process with the question: Do I want to share this information with the world, a selected few, or no one at all? Your response to this question will help you gauge the level of security you will need to implement in order to defend your data and its environment against internal and/or external unauthorized access.

2. What is TCP/IP?

In order to understand why a layered security approach is so effective, you first need to understand some TCP/IP basics. "TCP/IP is a protocol suite that allows computers of all sizes, from many different computer vendors, running totally different operating systems, to communicate with each other."[1] It was developed in the 1970s by two pioneering network engineers named Vinton Cerf and Bob Kahn. TCP/IP was later adopted by the Department of Defense to provide fault tolerant communication between U.S. military leaders in case of a Nuclear War. Its rules and standards primarily govern Internet communication today.

[1] Stevens

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

The TCP/IP Protocol Stack is made up of four primary layers: the Application, Transport, Network, and Link layers (Diagram 1).

Illustration of the TCP/IP Model

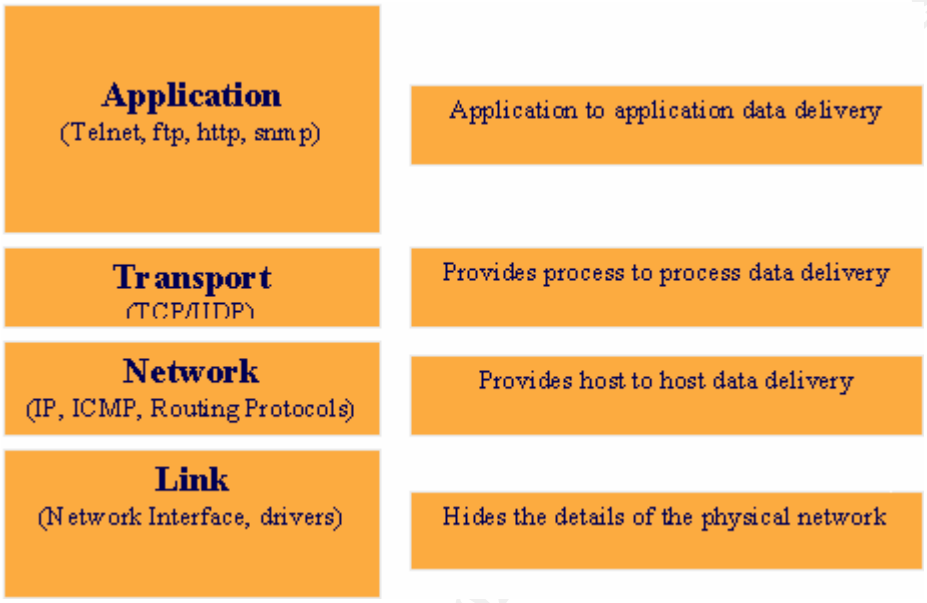


Diagram 1

Each layer within the TCP/IP protocol suite has a specific function. When the layers of the model are combined and transmitted, communication between systems can occur. Through properly understanding each layer's function, you can either set out to exploit or secure information packets transmitted from one system to another. Understanding how the layers interact will provide the basis for implementing security at the different layers. In order to illustrate the practical application of defense in depth concepts, I will provide a brief description of defense in depth; and for each layer of the TCP/IP model, discuss its function and provide a few examples of layered security.

3. What is Defense in Depth?

Defense in depth is a series of barriers, where each barrier represents a different challenge such that if one defense fails another defense is in place. Consider the case of polarized filters, which you may have had in science class. By taking the first polarized filter and using it to saturate the colors only the light waves on the same plane are able to pass. Add a second filter with the polarized screen rotated 90 degrees from the first and no light passes. The effect of applying multiple filters is the elimination of light. In this case light can be considered network traffic and the filters are your defense barriers. Through strategic implementation of these barriers, various levels of security can be achieved.

Applying such concepts can help prevent direct attacks against important systems and avert easy exploitation of your networks. In addition, a defense in

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

depth strategy provides natural areas for the implementation of intrusion detection technologies. Ideally it should supply you enough information to allow for quick detection and effective response to a breach thereby reducing its impact.

Illustration of Communication

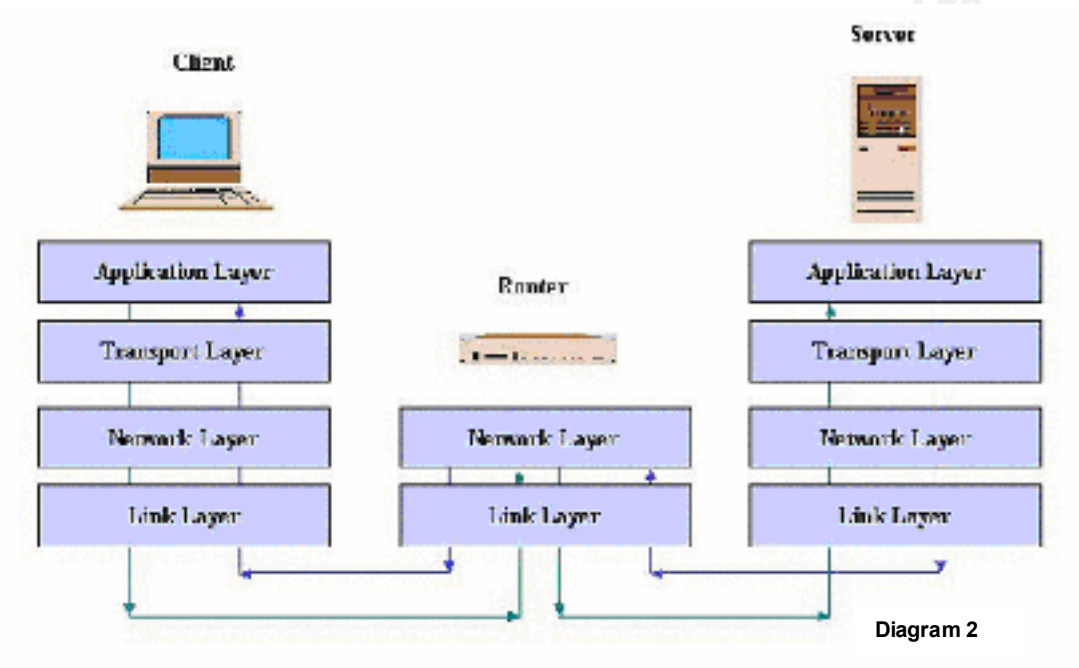


Diagram 2

4. Link Layer

The link layer is the lowest layer within the TCP/IP stack and its primary responsibility is to define how a computer connects to a network. A host can use a number of network topologies for communication such as Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI). The link layer allows the systems to identify the topology used and select the appropriate drivers or common language for communication. In other words the link layer does not regulate the network topology that the host is on but the network topology dictates the drivers that the link layer uses. By determining how communication is taking place, the link layer can determine how it will send and receive data.

Along with being the first layer of the TCP/IP model, the Link Layer is also the first opportunity to defend the network. Given the obvious reliance of wired LANs, anyone gaining access to the physical wire can damage or compromise the integrity and security of information on that network. Physical security is not only limited to wire based implementations. Wireless LAN's are just as susceptible to intrusion since walls and doors do not provide sufficient containment of the wireless signal. Without the proper security measures in place, even registered users of the network may be able to access information that would otherwise be restricted. Therefore, protecting and limiting authorization to all physical

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

components is the first line of defense. Examples of defense in depth at this layer can consist of physical security and Media Access Control (MAC) layer filtering.

4.1 Physical Security

Physical security is a vital part of any security plan and is fundamental to all security efforts. Physical security is not limited to unauthorized access; it also involves the protection of building sites and equipment (including information and software) from a number of harmful situations such as theft, vandalism, natural disaster, manmade catastrophes, and accidental damage. Implementing appropriate physical security controls requires solid building construction, suitable contingency plans, reliable power supplies, adequate climate control, and appropriate protection from intruders. Access to equipment should only be given to those that require access such as an administrator who provides support for your network.

4.2 Media Access Control (MAC) Layer Filtering

A MAC address is the numeric code that identifies a device on a network, such as your network interface card or cable modem. This hardware address is unique to all the network cards or cable modems ever manufactured. By copying the MAC address of the network device into the access table of a bridging device, you can ensure that only specified devices or computers can access your network through that bridge. A bridge is a device that connects two similar networks or divides one network into two sections. It works by reviewing the traffic frames that pass through it, determining which host on the network subnet it belongs to, and then relaying the traffic to that host. A bridge relays traffic by reviewing the source and destination address within the MAC layer information.

This type of filtering can be used as an additional barrier when implementing a wireless network. Most access points can transmit signals up to 500 or more feet, which in any direction will put the signal out onto another floor of your building, on the road or even out into the parking lot. A knowledgeable hacker, using a wireless card and tools that can be found on the Internet, can take advantage of receiving such a signal. In order to enhance the security within your network, the administrator will want to specify users and identify the MAC addresses for each wireless card they possess. Since, the MAC address is unique, only those wireless device addresses that you have specifically entered into the table will be able to communicate with the access point (Diagram 3). As you can see MAC layer filtering is a simple yet effective way to limit the level of access to your segmented network.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Illustration of MAC Address Filtering

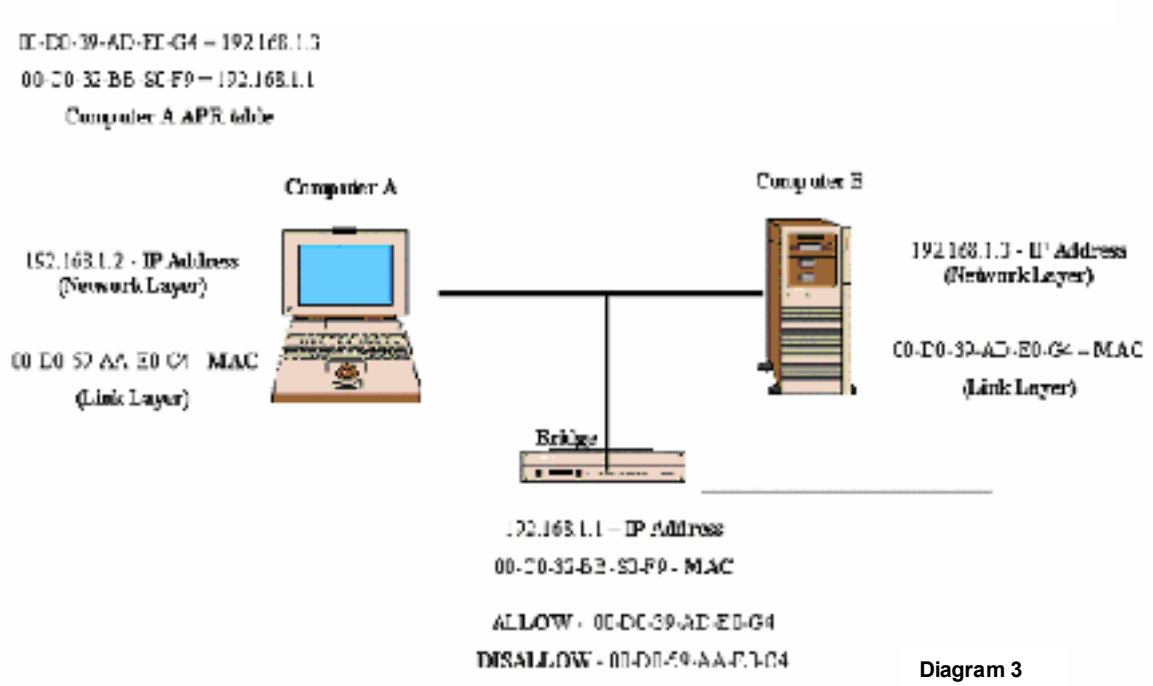


Diagram 3

5. Network Layer

The Network Layer is the second layer of the TCP/IP stack, which handles interaction with the network address scheme and connectivity over multiple network segments. It allows communication between systems by creating logical paths where the two hosts can communicate as though they have a dedicated connection, when in fact, packets might actually travel very different routes before arriving at their destination. The Network Layer ultimately defines how systems on different network segments find and communicate with each other using routers and IP addresses. An IP address is a unique logical address given to the host by a network administrator to identify a machine within a specific network. As information packets are passed down the TCP/IP stack, a source and target IP address are put into the IP header. Each IP address is made up of a host portion and a network portion. The host portion of an IP address refers to a particular system/device on a network while the network portion identifies the network segment where that system is located. Combined, they make up an IP address. Using the IP address, the IP protocol can determine whether the destination is local or remote by comparing the address to its own source address. When the destination IP address is not local, a router is used to transport the information over to its destination network and host. Routers basically move packets from one network to another as well as control traffic within its network segment.

Network security as it relates to the Internet is one of today's highest-profile information security issues. Companies and individuals alike need to weigh the costs and benefits of opening a connection between their private networks and

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

the unknown users and networks that compose the Internet. A number of attacks against the Internet Protocol are possible. These include attacks such as spoofing and session hijacking, which can lead to individuals revealing otherwise confidential information. Typically, spoofing and session hijacking exploit the fact that IP does not possess a robust mechanism for authentication. A packet can claim to originate from a given address when in fact it is originating from an unknown/untrusted source. Given the potential vulnerabilities of the network layer, it is important to consider ways of mitigating these risks. Increasing your network's security at the Network layer can consist of implementing a firewall, using Access Control Lists (ACL) for secure routing, and/or establishing a Virtual Private Network (VPN) to protect your data in transit.

5.1 Firewalls

A firewall is a logical or physical device that sits between the local area network (LAN) and the outside world that monitors and controls the flow of electronic traffic in both directions. It provides protection for user data and services against risks associated with Internet connectivity. When configured properly, authorized users are not even aware of its presence, but unauthorized users are blocked at the perimeter. This adds a layer of protection against general Internet risks. To be effective, a firewall must control all routes and services into and out of the network.

There are various types of firewalls that work at different layers of the TCP/IP stack. An example of a network layer firewall is the Cisco PIX. It offers stateful packet inspection and connection-oriented firewalling. Basic firewalls are considered to be stateless because they only look at individual data packets and do not keep track of connections already established; therefore every connection made is considered to be a new connection. However, A stateful firewall keeps track of more than just the information contained within a data packet. It attempts to keep track of the network connections at the application level in real time. As the firewall reviews the data packets flowing to an application, it classifies its properties and analyzes how it fits into the overall communication flow. As it identifies these properties it stores the data in a database, which it later uses to compare new data traffic. If a data packet does not fit into the general structure/classification for a particular application, that traffic is not allowed. An additional benefit to using a stateful firewall is that it can also be configured to redirect connections for authentication services and enhanced access control.

In short, an appropriately configured firewall such as a PIX, can serve as an effective layer of protection between your network and unauthorized individuals.

5.2 Access Control Lists

A router with an appropriately configured access control list (ACL) is a simple way to increase confidentiality within your network. Many routers now have the ability to selectively perform their functions based on a number of facts about a packet that are presented to it. In essence, ACLs serve as a packet filter, filtering out anything that does not agree with the list of accepted traffic. By default, a

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

router will normally pass all traffic sent to it with no restrictions. However, ACLs can be used to limit the types of packets that are allowed to come in and out of a network through the router in question. There are other benefits that can be gained by employing ACLs on a router such as increased network performance.

5.3 Virtual Private Network

A virtual private network (VPN) is a network that traverses a shared or public infrastructure, like the Internet, and allows you to establish a private and secure connection using specific protocols such as the Internet Protocol Security Standard (IPSec).

IPSec establishes secure, encrypted communication at the network level between firewalls, routers, and remote access devices. It also ensures: 1) authentication by validating the identities of communicating parties, 2) integrity by protecting data from alteration in route, and 3) privacy by safeguarding information from interception.

Ultimately, VPNs provide the ability for trusted parties (e.g. business partner, remote office/employee, etc.) to communicate with each other in such a way that it seems like they are directly connected over a private leased line. The session between them, although over the Internet, is both private (because the link is encrypted) and convenient (because it's through a public medium).

The Internet has transformed the way organizations and workers communicate by taking advantage of Internet infrastructure to extend access to internal network resources in a secure and cost effective manner. Information transmitted on the Internet is vulnerable to interception and tampering unless it is consciously protected using tools such as an IPSec based VPN. It is important to note that a number of firewall vendors such as Cisco and Checkpoint do include the ability to build VPNs either directly with their base product or as a separate purchase. If you have the need to connect several offices together, this may be an option.

6. Transport Layer

The third layer of the TCP/IP model is known as the Transport layer. The Transport Layer interacts with the application data and prepares it to be transmitted across a network. This layer ensures reliable connectivity, error recovery, and flow control from end-to-end through the sequencing of data packets during transmission using either TCP or the User Datagram Protocol (UDP) over IP.

6.1 Transmission Control Protocol (TCP)

TCP is connection oriented, which means that information is only sent once a connection is established. Through its functionality, TCP ensures that any information sent over the network is received by the other system making it very reliable in it's delivery. TCP achieves this by first establishing a connection via a three-way handshake and sending an acknowledgement back to the sending

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

host, while communication is established, to verify receipt of the data packets. TCP works well when all of the data is required to make sense of the information that is being sent.

6.2 User Datagram Protocol (UDP)

UDP on the other hand is a much simpler protocol and does not establish a connection before sending data. Therefore UDP is connectionless, meaning that it does not require the receiving host to acknowledge that it has received the information sent by the sending host. This makes UDP somewhat unreliable but it works well when TCP would be either too slow or where one or two missed packets won't make much of a difference to the end result such as with streaming video.

The Transport Layers protocols TCP and UDP both mainly focus on delivering information and not securing data while in transit. Since this layer only deals with these two protocols it is limited in what type of security can actually be implemented. Although, security at the layer itself is limited, there are certain add on security protocols, which interact with the Transport Layer in order to secure data while in transit. One such widely used protocol is Secure Socket Layers (SSL)

6.3 Secure Socket Layers (SSL)

SSL was developed by Netscape Corporation mainly to secure interactions between user and web servers; however, it is flexible enough to be used for other tasks as well. The SSL security protocol, now built into most major browsers, provides data encryption, server authentication, and message integrity features. Today, SSL is considered to be the industry-standard for protecting web communications and is widely used to transport highly sensitive information (e.g. credit card, loan applications, etc.) over the Internet. SSL comes in two strengths, 40-bit and 128-bit. The bit information refers to the length of the "session key" generated by every encrypted transaction. The longer the key, the more difficult it is to break the encryption code. Although SSL supports both 40-bit and 120-bit encryption, it is strongly recommended that one use the stronger 128-bit encryption as it provides much better security and it is no more difficult to implement than 40-bit encryption.

7. Application Layer

The final layer of the TCP/IP model is known as the Application Layer. The Application Layer is responsible for managing, setting up, coordinating, and terminating conversations between the applications at each end of the communication. It also makes sure that the data from the sending host is received in a format that is acceptable to the receiving host while supporting software application and end-user processes. In other words, the Application Layer interacts with the software application to determine whether network services will be required. For example, HTTP is an application that actually runs at the application layer while Internet Explorer or Netscape browser are software

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

applications that interact with the HTTP application. Additional examples of core Application Layer protocols are such applications as Telnet, Remote Login (rlogin), and File Transfer Protocol (FTP).

Since, TCP/IP on its own possesses minimal security functionality, application developers have taken it upon themselves to develop tools to ensure security at the Application Layer. Some examples of security tools at this layer are Application Proxy firewalls, Host based firewalls, and Anti Virus Software.

7.1 Application Proxy Firewalls

An application proxy firewall normally sits between an internal client and any externally held server. It is used to communicate on behalf of the application client by accepting the traffic initiated internally, identifying the source and port information, and then establishing a separate connection to the destination machine if allowed by the rules established. The internal client never communicates directly with the external server therefore the external server never has access to the internal network.

Proxy firewalls are very flexible in the way they can be configured. They can be used to allow or disallow connections and application functions based on a number of factors such as the user name, application used, protocol, source address, destination address, etc. Because this type of firewall is highly configurable, it can be leveraged in many different ways to further the depth of your network's security.

For example, your business may be required to submit a file to a business partner once a day for reporting purposes but does not want anyone internally to be able to download files outside of the internal network. In this case, the administrator may want to allow FTP PUT functionality to push the file out to the business partner's server but disallow FTP GET functionality to avoid downloading any virus-infected files. The administrator may want to take further precaution not only by limiting the application's functionality, but also by limiting the number of systems that can actually perform this task. He can also require that the user authenticate prior to allowing the outbound connection. All of this can be achieved using an application proxy firewall as long as it is accurately configured.

7.2 Host Based Firewalls

Previously we discussed the use of packet filtering and stateful inspection firewalls, which reside at the Network Layer. Similar to firewalls implemented at the Network Layer, there are also firewalls available that can be implemented at the Application Layer known as application level or personal firewalls. Examples of Application Layer Firewalls are Black Ice Defender, Norton Personal Firewall, and Zone Alarm.

Application level firewalls understand inbound and outbound session requests and work in real time to address threats before they reach the application. Firewalls at this layer of the TCP/IP model do not route traffic on the network,

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

instead all traffic stops at the firewall and the firewall initiates its own connections based on configured rule sets. As the requests come to the application layer, the firewall parses the traffic and associates it with an already established session or creates a new session if necessary. Once the session is either matched or established the firewall will compare requests to the session policy created by the user/administrator. Since most application level firewalls have the ability to prompt the user to allow or reject network connections, the best way to initially implement a personal firewall is to deny all traffic unless specifically granted by the user. Once the user makes a decision the firewall will record it as a rule. If you are less familiar with using personal firewalls, an automatic configuration setting can also be selected making configuration and setup quite simple. Application level firewalls can be easily implemented and provide a decent level of security at the individual host level.

7.3 Anti Virus Scanner

Patching security holes as they develop is good practice, but a good security solution must always provide protection from files that come into the network via e-mail, Internet downloads, or floppy disk. Anti Virus software, although certainly not fool proof, can provide an additional level of protection for any host against known virus threats. The software works by detecting specific traits (signatures) of well-known and documented viruses. If while scanning, the Anti Virus software comes across a known signature, it will automatically alert the user that a virus has been detected. For example, Symantec's Norton Anti Virus "automatically checks for newly discovered threats and periodically scans systems for those threats. It also watches in real time while new files are downloaded from the Internet or detached from e-mail messages to make sure nothing unsafe gets through."[2]

8. Conclusion

Given the explosive growth of the Internet and its increasing business dependency, it is only a matter of time until your business also, if not already, has a presence on the World Wide Web. Any business that wants to remain competitive will need to be connected. However, connecting your network infrastructure to the Internet without the proper security controls is simply not a good idea. Doing so would be like jumping into the ocean with 100-foot waves; you may tread water for a while but eventually you'll be swallowed whole. Opening your network up to the Internet will inevitably invite attempts at probing into your network and you must be prepared to defend it. Keep in mind that each day it becomes easier for malicious individuals to steal, change, or simply delete important information; the tools they require are widely available on the Internet and the exploitable vulnerabilities are numerous.

Unfortunately, there are no silver bullets or single implementations that can address all network security concerns. In order to maintain the confidentiality,

[2] Symantec

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

integrity, and availability of your information, a layered security defense that marries both policy and tools at every layer of the communication model is the only secure alternative.

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

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

9. References

9.1 Books/Guides

W. Richard Stevens. *TCP/IP Illustrated vol. 1 The Protocols*, New York: Addison Wesley 1994.

9.2 Papers/Links

The Moschovitis Group. "History of the Internet", 1999.
http://www.historyoftheinternet.com/chap4.html, June 2002

Hedrick, Charles L. " Introduction to the Internet Protocols", 1987
http://www.doc.ic.ac.uk/~ih/doc/pc_conn/tcpip/intro/intro0.html, August 2002.

Ferrell, John. "Guide to Networking for K-12 Schools", February 1998
http://www.netc.org/network_guide/, July 2002.

Cisco Systems. "Data Link Switching", February 2002
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/dlsw.htm#xtocid1, August 2002.

"Master a Network part I", August 2001
http://www.pchardware.ro/Articles/article.php?id=124, July 2002.

Brooke, Paul. "Building an In-Depth Defense", July 9, 2001
http://www.networkcomputing.com/1214/1214ws1.html, July 2002.

Whipple, William L. "TCP/IP for Internet Administrators" 1997
http://www.pku.edu.cn/academic/research/computer-center/tc/FrameMain.html, June 2002.

Cisco Systems. "Cisco IOS Lock and Key Security", July 2002
http://www.cisco.com/warp/public/cc/pd/iosw/ioft/iolk/tech/landk_wp.htm, July 2002.

Symantec. "Defense in Depth Benefits"
http://securityresponse.symantec.com/avcenter/security/Content/security.articles/defense.in.depth.html, August 2002.

Cisco Systems. "Cisco PIX Firewall", September 2002
http://www.tribecaexpress.com/ciscoPIX.htm, August 2002.

Netscape Communication Corp. "How SSL encryption works", 1999
http://support.acmeinternet.com/howtofaqs/ecommerce/ssl-howitworks.htm, June 2002.

Hamann, E. M. "Securing e-business applications using smart cards", November 2001
http://researchweb.watson.ibm.com/journal/sj/403/hamann.html, July 2002.

Stiller Research. "How to get the most from your Anti Virus Product", August 1998
http://www.stiller.com/avsw.htm , June 2002.

Patel, Rahul & Friend, Robert. "Taking a Stateful Approach to Firewall Design", April 2002,
http://www.commsdesign.com/design_corner/OEG20020404S0030, September 2002.

Steinke, Steve, "Firewalls", June 2000
http://www.networkmagazine.com/article/NMG20000613S0010, July 2002.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.