

Internet Security Primer

July 2000

Executive Summary

It is simply not possible to achieve absolute network security. Network system managers must constantly strive to attain absolute network security because networks are vital bridges of communication through which users carry out crucial transactions and dialogue. Since security and privacy are the antithesis of sharing and distribution, network security must be a balance of the two. Managers must dually facilitate the risks of losing information and the provision of security to prevent unwanted intrusion. This primer is intended to help management successfully administer network security initiatives by providing an overview of security principles and technologies that are appropriate in network security today and into the future.

Introduction

Secure environments, by definition, are not conducive to networking. Security may be achieved when the information is isolated, locked in a safe, surrounded by armed guards, barbed wire fences, and placed in Fort Knox. Even these measures do not ensure absolute security in the face of a well-planned attack. It is simply impossible, therefore, to render a network system completely secure, and any reader who wishes to understand and apply the principles of security to the Internet or any other network must first accept this basic tenet. Nevertheless, every company with internet, intranet, and extranet embedded infrastructures should have the tools and knowledge to strive for a perfect balance between optimal information flow and industry leading security.

Advances to remedy the problem of data isolation create exponential rises in data vulnerability. "Islands of Automation" were a hindrance to conducting business successfully, as critical information required by one "island" could not be accessed in a practical and reliable manner by another "island". Networks were created as a remedy to the problem of data isolation in the early days of computing.

It wasn't until 1969 that the US Dept. of Defense, and its Advanced Research Project Agency, began to experiment with a computer simulation of a message routing scheme. Its purpose was to replace switchboards that relied on human operators to provide dedicated communication links. Interconnecting nodes/links permitted a network to suffer a large number of breaks, by reconstituting itself through other nodes (points on the net) by rapidly relearning to make best use of the surviving links. Networks became communication bridges by which "islands" could be integrated. Yet, "islands" of critical information were then exposed in exponentially greater numbers of data thoroughfares.

Since security and privacy are the antithesis of the sharing and distribution processes encouraged by the internet, network security managers must walk a thin line between providing appropriate access to those who need the information and safeguarding it from unauthorized entities. Network security managers must determine which level of risk is appropriate to assume according to the sensitivity of the information being guarded.

The explosion of networks across the United States, and the rest of the world, has raised the specter of corporate espionage to new heights. Today's corporations understand that substantial assets can be lost when networks are compromised. Organizations that retain control of their information can gain a substantial competitive advantage: those who do not are vulnerable to losing valuable trade secrets to competitive spies and malicious attacks.

Intra-company acts of sabotage or ignorance can be just as costly as outside attacks. New means of interconnectivity bring new challenges to corporate security. As more business is done through mobile communications, employees become increasingly capable of compromising the information that they demand while on the road. Faced with networks of increasing size, power, and speed, managers must reconcile the risks of losing information critical to the enterprise's operation with the costs and constraints associated with an overly aggressive security solution. This primer is intended to help management successfully balance these variables by

providing an overview of security principles and the technologies that are appropriate for securing digital networks.

Network Security Issues

Basic Security Concepts

A good place to begin is by defining the basic concepts involved in securing any object. Four key words in any security lexicon are vulnerability, threat, attack, and countermeasure. An examination of each follows.

Vulnerability is the susceptibility of a situation to being compromised. It suggests potentiality of weakness through existing holes or deficiencies. Vulnerability in and of itself may or may not pose a serious problem, depending on what tools are available to exploit that weakness. The classic analogy of vulnerability comes to us from Greek Mythology, with the story of Achilles, whose heel represented his greatest vulnerability.

A threat is an action or tool that exploits and exposes vulnerability, therefore compromising the integrity of a network. Not all threats are equal in terms of their ability to expose and exploit vulnerabilities. For example, the Microsoft Concept virus exploits vulnerability in Word Macros allowing access to the users' file system, but the virus itself is relatively benign. With slight modifications, this virus could do much more damage.

An attack represents the method in which a particular threat could be used to exploit vulnerability. It is entirely possible that situations could exist where vulnerabilities are known and threats are developed, but no reasonable attack can be conceived to use the specific threat upon a vulnerability of the system. An example of an attack is a Trojan Horse attack, where a destructive tool such as a virus is packaged within a seemingly desirable object, like a piece of free software. The success of the attack is contingent upon the stealth of its vehicle or Trojan Horse.

Countermeasures are those actions taken to protect systems from. Achilles covered his heel with a protective metal plate as a countermeasure to potential attacks to his one vulnerability. In the network security world, countermeasures consist of tools such as virus detection and cleansing, packet filtering, password authentication, encryption, and firewalls.

Any security scheme must identify vulnerabilities and threats, anticipate potential attacks, assess whether they are likely to succeed or not, assess what the potential damage might be from successful attacks, and then implement countermeasures against those defined attacks which are deemed to be significant enough to counter. Therefore, we can see that security is all about identifying and managing risk, and that security is a very relative concept which must be tailored to the needs, budget, and culture of each organization.

For example, a Trojan Horse attack on one organization could succeed in compromising extremely important information. The same attack on another organization might only result in minimal damage, perhaps because there is no sensitive data available on the victimized system. Furthermore, companies have personalities just as people do, and therefore, some companies are willing to live with more risk than others. In each of these organizations, different security

schemes should be employed with different countermeasures to suit their specific situations.

As we will discuss later, management must consider all of these factors in defining a security strategy. Management must also consider the cost of protecting against all possible attacks. Security is a significant investment, and each organization must determine how much it is willing to invest in appropriate countermeasures. Once the bottom line is determined, organizations determine can focus their choice of defenses.

Generic Security Threats

In any organization, there are a number of generic security threats that must be dealt with. These include the theft of information, the compromising or corruption of information, loss of confidentiality, and the disruption of service.

One of the major threats companies are dealing with is the introduction of malicious programs over the network. The term "computer virus" has been used loosely to categorize these attacks come in Trojan Horses, worms, and logic bombs as well as true viruses.

A Trojan Horse is a program that conceals harmful code. It usually disguises itself as an attractive or useful program lures users to execute it, and in doing so, damages the user's system. For example, a posting in the US Department of Energy Computer Advisory Capability page lists a known Trojan Horse in a program called AOL4FREE. While the title suggests that this program will allow you to participate in AOL without any costs, running the program will delete all of the files on your hard drive. The program hidden in the Trojan Horse can be one that causes malicious damage, or one that performs some espionage for the attacker, such as stealing the password file from the computer it invades.

A logic bomb is code that checks for certain conditions and when these conditions are met, it "detonates" to do its damage. Sometimes, like the Magellan virus, the trigger logic is a date, but it can be any given set of parameters, including a person's name, a bank account number, or some combination of events and parameters.

A worm is a self-contained program that replicates itself across the network and therefore multiplies its damage by infecting many different nodes.

A virus is code that plants a version of itself in any program it can modify. The Microsoft Concept virus is a good example: once it has "infected" Microsoft Word, all subsequent documents which are opened by the user may only be saved as template files. In all other respects, Microsoft Word continues to operate normally.

It should be noted that these are not mutually exclusive threats. A logic bomb could plant a virus under the specified conditions, as could a Trojan Horse deliver a worm.

Furthermore, each of these threats could have single or multiple missions, such as the theft of data, the compromising of confidentiality, integrity or availability, or the disruption of service to the organization. Threats also encompass the theft or compromise of information while it is in transit between endpoints of a network. One such example is called Snooping, in which an attacker simply eavesdrops on electronic communications.

These are the classes of threats that today's network managers must deal with, and that senior management must be aware of, since they will play a major part in determining the appropriate and tolerable cost of security to counteract these potential threats.

Security Countermeasures

Given the above scenario, a reasonable question at this point might be: "What tools are available today to help mitigate the consequences of these security threats on the network?" The good news is that there are multiple technologies that can be brought to bear on the issues, and they are impressively effective.

Security Policy

The bad news is that no amount of technology can overcome a poorly planned, poorly implemented or nonexistent security policy. Consider the following story witnessed by a poster on a security newsgroup on the Internet: A customer being waited on at a public service agency (say a Department of Motor Vehicles) requires some information from the clerk who in turn needs to access that information from a workstation centrally located in the area behind the window. Sitting at the workstation, the clerk yells to a co-worker "Mary, is the password still ----- ?"

Again, the security of any information in any organization today is primarily dependent on the quality of the security policy and the processes by which that organization imposes on itself. If the security procedures are lax, are not enforced uniformly, and allow gaping security holes to exist, no amount of technology will restore the security breaches. Organizations that are concerned about security on the Internet should ask themselves a few of the following questions before worrying about encryption, packet filtering, proxy servers, and other related technology solutions.

- ?? Does the corporate policy allow passwords such as
 - o "password",
 - o employee initials, or
 - o names or initials of employees' immediate family members?
- ?? Is there a process in place to change passwords periodically?
- ?? Do employees keep their password written on paper under their mouse-pads, keyboards, monitors, etc.?
- ?? Is there a program in place to make employees aware of the need for security and to disseminate security procedures to the employees to facilitate its implementation?
- ?? Do employees understand the different levels of security of information and what techniques to apply to each to ensure an appropriate level of protection?

- ?? Is the responsibility for information security assigned to a senior member of the management team?
- ?? Is there a set of guidelines to identify the security classification of different documents and information generated by the employees; is there a process in place to classify or categorize these documents and information and secure them appropriately?

It should be self evident that the primary need for any organization is to get its own house in order, identify its security needs based on the types of information with which it deals and develop a security policy and plan before committing to technology. Following are some elements of a good security plan.

- ?? Develop security requirements based on an analysis of the organization's mission, the information at risk, the threats to that information and the implications of any successful attacks.
- ?? Appoint a security officer and delineate clearly the required job responsibilities and skills.
- ?? Define appropriate security services and mechanisms and allocate them to components of the company's IT systems.
- ?? Identify different measures of security appropriate for each level.
- ?? Identify users who should have access to each level of security.
- ?? Remember that security is not only technology; physical security and procedural security are as important as the technology used.

Authentication

A primary tool in securing any computer system is the ability to recognize and verify the identity of users. This security feature is known as authentication. Traditionally, special names and secret passwords have been used to authenticate users, but as the anecdote above demonstrates, the password is only as good as the users' ability to keep it secret and protect it from being abused by unauthorized users.

There are three generally accepted techniques for authenticating users to host machines.

1. Authentication by something the user knows

This is the password/username concept described above. There are two common approaches to password authentication, known as PAP and CHAP. PAP, which stands for Password Authentication Protocol, simply asks the requester to provide a "secret" password, and if the password provided is included in the user profiles, the requester is given access. CHAP (Challenge Handshake Authentication Protocol) takes the concept one step further by challenging the requester to encrypt a response to the challenge message. This, in effect, acts as a different password for each entry. Often, the CHAP mechanism is combined with an encrypting smart card, which uses

an encryption key to encode the challenge message. Only if the challenge message is correct will the requester be granted access to the system.

2. *Authentication by something the user has.*

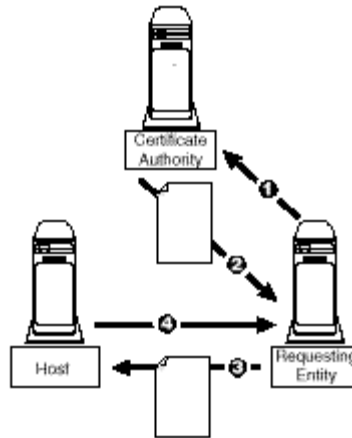
In this technique, the user is given some kind of token, such as a magnetic stripe card, key, or in sophisticated cases such as the remote access standard RADIUS discussed later, the user has a smart card equipped with a computer chip which can generate an encrypted code back to the computer system.

3. *Authentication by physical characteristics.*

Here, the mechanism is to recognize some measure of the individual that ostensibly cannot be duplicated. Biometric techniques such as fingerprint ID, palm print ID, retinal scan, manual and digital signature, or voice recognition are used to validate the identity of the potential user.

Authentication is also necessary when two computers communicate with each other. For example, what should my host computer do when another computer asks to have a disk mounted containing all of my organization's personnel data? How do I know that the requesting computer has a legitimate reason to access that information, and that it is not an external network hacker trying to steal information from my organization? In order to prevent such events, the Internet Engineering Task Force (IETF) has formed a working group by the name of IPsec (Internet Protocol Security). Additionally, there are a number of *de facto* standards - those which are developed by companies rather than by official committees, but which enjoy widespread acceptance. While many of these standards are under review and take some time to work their way through the approval process, two are worthy of mention here, IPsec's SKIP (Simple Key Management for Internet Protocol) and Livingston's RADIUS (Remote Authentication Dial In User Service).

SKIP is a technique for providing authentication and encryption security at the IP layer of the Internet architecture. It relies on the existence of an authority in the network that issues certificates to known and trusted entities within the system. If an entity claiming to be a member of the system requests an action, the receiving computer system can have the requester present an encrypted certification that they are who they say they are. The certificate conforms to one of the methods of authentication, namely, a secret encoding technique and a secret key, which are only available to trusted members of the system. The fact that SKIP operates at the very lowest protocol layers of the architecture has the advantage that it will protect all upstream applications as well, by preventing connections between systems which are not authorized. Since potential intruders cannot even establish connections, their ability to do malicious damage is severely restricted.



In the scenario depicted above, the Requesting Entity first gains a certificate by requesting one from the trusted certification authority (1), who validates the trustworthiness of the Requester by granting a certificate (2). Armed with this certificate, the Requester can now petition the host and present the certificate along with the request of the host (3). The host, upon seeing the certificate will grant the information to the requester (4).

RADIUS is one of the more popular public network authentication protocols. The primary purpose of RADIUS is to offer centralized access control for remote dial-in users. RADIUS simplifies the administration of passwords, user names, profiles for remote users, and other security and accounting related information by placing all of the security in a central server, and issuing challenges to the user.

Virtual Private Networks

The Internet Community is constantly seeking new and better mechanisms to secure the Internet. Today, there are several other relevant proposals for standards, which are under review by the Internet Engineering Task Force (IETF). One generating some potential interest is the Level 2 Tunneling Protocol (L2TP), which is under review as part of the IPsec group within the IETF. This proposal would establish a set of protocols by which compliant internet components could create their own channel *inside* the Internet. This channel would be protected by authentication and encryption countermeasures. These would ensure that even though the traffic is being transmitted over the public internet, individual sessions can be established which are private to those members allowed to work within that channel. The technology is known as *tunneling* because the correspondents are creating a tunnel of sorts through the public packets inhabiting the internet, and exchanging very private communications within them. The concept comes from medieval times, where tunnels were built between fortified towns and castles to allow their inhabitants to move safely between them away from the dangers of the bands of marauders outside their gates.

The use of tunneling technology allows another concept to be implemented: the concept of a Virtual Private Network, or VPN. Companies who want a less expensive alternative to private Wide Area Networks can utilize tunneling within the Internet and develop their own virtual WANs, safe from unwanted intrusions, yet riding on the cost benefits of the Internet mass volumes.

Non-Repudiation

This security concept protects against the sender or receiver denying that they sent or received certain communications. For example, when a person sends a certified or registered letter via the United States Postal Service (USPS), the recipient is supposed to prove his or her identity to the delivery person, and then confirm their receipt by signing a form. The signed form is then returned to the sender, which proves to the sender that their correspondence was delivered. This prevents the recipient (for example a debtor) from claiming that they never received the correspondence (for example a demand note) and therefore using that as an excuse for their actions (not paying the debt). In computer networks, these kinds of services are also available, and are becoming increasingly valuable as commerce on the Internet continues to gain in popularity.

There are three different types of non-repudiation services that are applicable in computer network messaging:

- ?? non-repudiation of Delivery Service,
- ?? non-repudiation of Origin Service, and
- ?? non-repudiation of Submission Service.

Non-repudiation of Delivery Service is similar to the US Post office certified mail example above. This provides the sender with proof that a message was successfully delivered to the intended recipient. Many e-mail packages offer senders the option to request a return receipt. This return receipt provides the sender with a non-repudiation of delivery service feature - the recipient can't legitimately claim they did not receive the message.

Non-repudiation of Origin of Service provides the recipient with proof of who originated the message and what it contains. For example, according to a Usenet posting, America Online (AOL) was victimized by crackers pretending to be AOL employees requesting passwords and credit card information from subscribers. A recent article in the Los Angeles Times stated that a particular cracker "armed with an AOL hacker program created... [a] fake screen to pass himself off as an AOL employee and steal Knaiger's [the AOL user] password". Non Repudiation of Origin of Service could have foiled this kind of attack if it had been available to AOL subscribers. If it had been available, users could have verified that the crackers were not genuinely AOL employees, and therefore would not have given away their passwords.

Non-repudiation of Submission Service is similar to the concept of non-repudiation of delivery. This service offers proof that a given message was in fact sent from a particular sender. If we go back to the US Post Office example, when we mail important papers such as legal documents, it is considered prudent to send them via registered mail. When we do so, we get a receipt from the Postal Service and a special identification number is affixed to the return. Thus, if the recipient does not receive the documents, or contends that it was not sent on time, we have evidence that our submission did occur at a particular time.

Integrity

Integrity refers to the completeness and fidelity of the message as it passes through the network. The key here is making sure that the data passes from the source to the destination without undetected alteration. Note the use of the word "undetected". We may not be able to thwart someone from tapping out messages and attempting to modify them as they move through the network, but we will be able to detect any attempt at modification and therefore reject the message if such a modification attempt is detectable.

If the order of transmitted data also is ensured, the service is termed connection-oriented integrity. The term anti-replay refers to a minimal form of connection-oriented integrity designed to detect and reject duplicated or very old data units.

Confidentiality

Confidentiality is a security property that ensures that data is disclosed only to those authorized to use it. The key point behind ensuring the confidentiality of information on a network is to deny information to anyone who is not specifically authorized to see it or use it. Encryption is a frequently used mechanism for guaranteeing confidentiality, since only those recipients who have access to the decrypting key are able to decode the messages. Confidentiality therefore predicates privacy.

Access Control

Access control relates to the acceptance or rejection of a particular request to have access to some service or data in any given system. A service could be programs, and devices such as printers or file systems. Data encompasses text files, images, a collection of files, or any combination thereof. The most important question to be asked at this point is what are the risks involved in allowing access to any of the system's services or information to individuals requesting such access?

It is therefore necessary to define a set of access rights, privileges, and authorizations, and assign these to appropriate people within the domain of the system under analysis. Web advertisements necessitate mass access in order to be cost effective. Data collected from such advertisement web pages, highlighting traffic characteristics, necessitates strict access privileges.

Internet Architecture: An Overview

In order to understand better how these security principles can be applied, we need to understand the standard networking architecture and how the specific Internet Architecture fits this model. Then we can see how the security principles we have just discussed apply to that Internet model.

ISO 7 layer model


The International Standards Organization (ISO) published a groundbreaking network architecture design in the early 1980s. Its primary philosophy is that different standard and open "layers" of the architecture should handle different telecommunications functions. This so called Open Systems Interconnect (OSI) model is constructed as follows.

The very lowest layer is the physical layer, which is responsible for the physical transmission of data from computer to network. Here, there are electronic circuits and mechanical connectors that define how transmissions are to occur over coaxial ethernet, modems, FDDI or any other medium for transmitting data.

Next is the Data-link layer, which is responsible for the integrity of the bit stream between any two points. Here, there are standards for redundancy checks, parity, retransmission protocols, etc. to ensure that the same sequence of bits sent from point A is received at point B.

The Network layer extends the concepts of the Link layer into multiple networks, which may or may not be compatible. Internetworking implies that this layer must be aware of different routes available to connect the sender with the recipient.

The Transport Layer ensures that different transmissions, which may be part of a sequence and which may have traversed the network via different paths, are appropriately re-sequenced at the receiver's site.



The Session Layer manages the connecting and disconnecting of interactions between two computers and how the data is to be exchanged (duplex, simplex, etc.)

Presentation determines what code sets will be used (ASCII, EBCDIC, international character sets, etc.).

Finally, we come to the Applications Layer in which specific applications like FTP, Telnet, e-mail, Archie, and others reside.

The architecture of the OSI model is such that each layer uses services "below" it and provides services to those layers "above" it, giving the appearance of a stack. In fact, the model is known as a protocol stack and other architectures, such as TCP/IP will also follow the stack model.

Internet layers

At this point, the vast majority of people interested in the Internet are familiar with the acronym TCP/IP (Transmission Control Protocol/Internet Protocol), which form the foundation of communications for the net. This section will cover the architectural constructs of the TCP/IP structure and their relationship to the Open Systems Interconnect model.

Packet Switched Networks

The Internet uses the concept of packets and packet switching to allow for simultaneous access by millions of people. Transmissions between any two points are broken into smaller transmissions known as packets. By doing this, everyone's messages can be sent in an interleaved fashion so that all users see nearly the same level of performance.

Each connection to the Internet is designated with a unique Internet Address, usually written as four numbers determined by a standard *Internet Protocol (IP)*. Packets are shipped to Internet destinations through routers, which use Transmission Control Protocol, or TCP.


These fundamental layers of the Internet form the backbone upon which data can be sent from one point to another, with integrity.

Applications Layer

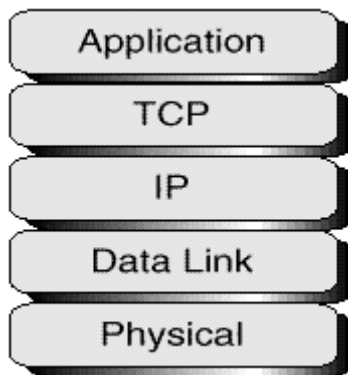
Getting data from point A to point B is essential, but it is not sufficient. It is the equivalent of finding a telephone number in a phone book and being able to establish a telephone connection between the US and a foreign country. The individuals will not be able to communicate unless there is a common language. Similarly, on the Internet, as specified in the OSI model, there are other protocols and tools which are specific to the application layer and which correspond to these telephone analogies.

- ?? File Transfer Protocol (FTP) is the oldest commonly used method to allow files transfer from one location to another. Using an anonymous FTP, a user can browse the files of a host, select appropriate files and have those files transferred to his or her own system.
- ?? Telnet is an application and set of protocols that allows a user's Internet connected terminal to act as if it were directly connected to the host computer.

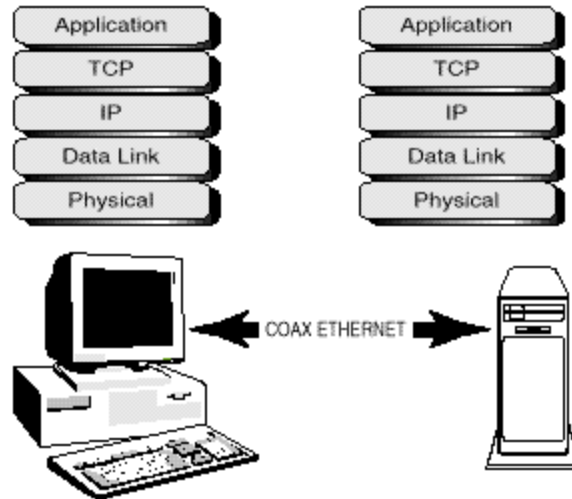
- ?? E-mail applications use the Internet as an interconnecting infrastructure to allow messages to be transmitted among users. For example, AOL and CompuServe use the Internet to move e-mail traffic to and from each other.



We can see, therefore, that the Internet layers roughly correspond to the OSI layers, with the exception of session and presentation. This implies that Internet applications must handle the tasks generally assigned to session and presentation in the OSI model.



To demonstrate how these layers work, the diagram below shows a PC connected to a Web server over a local Ethernet connection.



In this environment, the PC has a coaxial connection for the physical layer, and uses Ethernet as the link level control. The PC also will have a TCP/IP software protocol stack, and will likely use Netscape Navigator or Microsoft Explorer as its application software. On the server end, all of the bottom layers are identical, and the application will be the Web server and any custom software written, such as shopping or search engine applications. As mentioned earlier, each layer has its own set of dialogs and its own language in which to conduct those. The application layer, for example, the language and protocol are contained within HTML commands and responses.

Internet Security Architecture

Given the architectures described earlier, where are the vulnerabilities, and specifically, what countermeasures can be taken to thwart potential attacks in this architecture? This section will examine the security tools available, and their intended deployment, for maximum protection. The placement of these security components will constitute the Security Architecture for network configurations. Note that, as in any architecture, the components are designed with a great deal of flexibility and depending on particular needs of specific situations, the selection of components and their interrelationships may vary significantly.

Two Approaches to Security

Over time, two distinct approaches have evolved to applying security countermeasures: *networked coupled security* and *application coupled security*. As the names imply, the first philosophy favors the use of securing the network infrastructure, while the second builds security into the applications themselves.

Network Coupled Security

In a Network coupled scheme, the focus is to make the network itself a trusted and secure subsystem so that the applications can assume the data being transmitted is safe, comes from authorized users and is being delivered to the appropriate recipients. If the network itself is secure, then the applications don't have to do anything special to operate in a secure environment - they simply assume that all security functions being performed by the network itself. This

philosophy is very similar to that of the Internet and OSI architectures. Applications that operate in the OSI and Internet environments do not concern themselves with sequencing of packets, validation of IP addresses, etc. The applications assume that the layers below in the supporting protocol stacks have done their job, and therefore the applications can concentrate on the data content. Similarly, in a secure network environment, applications can assume that the security being handled by the lower levels.

The most significant advantage to network coupled security is that applications do not have to be security aware. Applications that are not security aware can be moved into a secure environment without modification. Less obvious, but equally important, is that the use of a consistent security mechanism within the network allows applications to interoperate from a security standpoint. There is no possibility that different applications will insist on different authentication schemes, key management schemes, etc.

Application Coupled Security

Proponents of this scheme argue that the application knows best what kind of security is required for that application. Therefore, control of the security aspects should rest in the application layer. For these proponents, the need to create security aware applications is not a disadvantage, but rather a natural and reasonable consequence of the need to apply security at that level. Similarly, application-coupling advocates tout their advantage of flexibility, arguing that a "one size fits all" approach in network security is insufficient for the broad range of security requirements.

Different Tools for Different Layers

There is no shortage of technology available to secure an organization's Internet connections. An organization should focus on which tools should be utilized and in what way.



Early on, router manufacturers recognized the key role they could play in this endeavor, and have placed filtering capabilities in their products to establish a primary front line of defense. Routers capable of examining and screening network traffic based on the IP packet addresses are known as a "screening router". Some advanced routers provide the capability to screen packets based

upon other criteria such as the type of protocol (http, ftp, udp), the source address, and the destination address fields for a particular type of protocol. This way, a communications manager can build "profiles" of users who are allowed access to different applications based on the protocols. Such a case is shown above.

In this same figure, the packet filtering of the screening router is enhanced with authentication software that can add either password authentication or challenge authentication. In the scenario described above, simple filtering, even with profiles, cannot authenticate that the individual on the other end of the connection is in fact the individual who should have access to the applications and data residing on the server connected to the local LAN. Therefore, we need to add PAP/CHAP authentication to accomplish this, which provides another layer of security to the system.

The screening router either alone or in combination with authentication, is known as a "Firewall", because it keeps the "fire" of unsecured communications outside a protective "wall".

Another popular configuration, particularly for organizations which have WWW presence, is the use of a so called "bastion host". Again borrowing from the medieval fortifications of cities, this concept consists of a double walled security layer. The outside wall, or perimeter wall, consists of a screening router which provides a first pass screen for the population of outside users who are allowed access to the Internet accessible applications in the Bastion Hosts, which sits in the "moat" between the two walls. A secondary router with or without authentication enhancements provides a second filter for those few privileged users who have access to the internal network. The bastion host concept is demonstrated below.

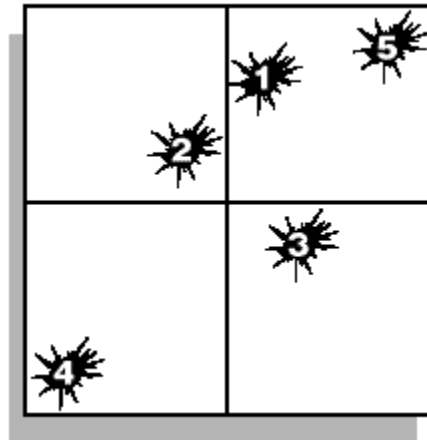


Protecting against unauthorized access to data, via controlling who is allowed to communicate with protected servers, will guard against many of the vulnerabilities of networks. However, such schemes do not protect from espionage and theft of data, which may be captured en route between two validated correspondents. It is in this area where the addition of encryption and key management, as defined in the SKIP standard, will provide effective countermeasures. By encrypting the data and properly managing the keys to the encryption and decryption, data which is intercepted is rendered unusable, and at the same time un-modifiable, thereby adding a further layer of protection for the data.

Beyond this level, additional security is still available when it's coupled to applications as mentioned earlier. For example, database systems also have authentication capabilities with user names and passwords as well as profiles, access control lists, and the like. It is possible to add yet more layers of security beyond those discussed so far by adding similar technologies to the application layer. In these schemes, applications could also issue challenge passwords beyond those required to gain access to the network, thereby increasing the security of the data by decreasing the odds that a single error (lost password, etc.) could compromise the application or the data. One common form of application level security is the use of Secure Socket Layer (SSL) directly coupled with the application. In order for SSL to work, both the Browser client and the Server application must support its use, making the application security aware. SSL is discussed further in the tools section of this paper.

Managing the Risk

Network security is all about managing risks and using this risk management analysis to provide appropriate security at an affordable price. This section will explore a Risk Management tool that can be used to analyze the risks in your organization and take appropriate countermeasures.



Risk Determination

Risks can be characterized by two criteria: the likelihood that a particular attack will be successful, and the consequences of the results if the attack is successful. Security measures, like any other operation, cost money and should be subjected to cost benefit analysis and risk assessment. Risk mitigation strategies focus on either minimizing the likelihood of occurrence (by employing countermeasures), or by minimizing the consequences of the attack. One way to depict this is to characterize the risks along two axes, one indicating increasing likelihood of an attack succeeding, and a second indicating increasingly dire consequences. The individual attacks are plotted according to the two axes, and depending on where they fall, they can be characterized as worthy of defending or not. If an attack is considered serious enough to defend against, countermeasures are developed to reposition the attack into the lower left hand quadrant.

In the previous chart, five potential attacks are plotted in a hypothetical scenario to demonstrate this technique. Attacks numbered one and five fall into the upper right hand corner, which means the analysis has shown them to be likely to succeed, and that the consequences are serious. Attack number three is somewhat likely to succeed, and the results appear to be moderately serious. Attack number four is not likely to succeed, and even if it did, the results would not be particularly damaging. Attack number two is also not likely to succeed, but if it did, the results would be damaging. Based on this analysis, an organization might opt to do nothing about attack number four, would definitely provide defenses against attacks one and five, and would optionally defend against number two and number three, depending on budget constraints.

Internet Security Toolkit

This section will describe the different technology tools available to deploy in development of a security architecture for any given network.

Secure Sockets Layer (SSL)

The SSL is inserted between the TCP protocol and the application protocol. The SSL protocol operates in two phases. In the first phase, the sender and receiver agree on the read and write keys to be used, and then in the second phase data is encrypted using the keys chosen in the first phase. Authentication and secure key exchange is also achieved using the RSA public key encryption algorithm.

HyperText Transport Protocol Secure (HTTPS)

By layering a version of SSL between HTTP and TCP, Netscape has developed a secure version of HTTP as well.

Proxy servers

A proxy server is a firewall implemented in a hardware unit such as a workstation on a NT server, rather than in a router. This device looks at all of the data in each packet, not just addresses and headers. In most cases, the proxy examines the content and replaces the network address in the packet with proxy destinations that are known to be secure. Besides hiding the network from the outside world, they provide more control over the actual data at the application level. However, because proxy servers inspect all of the data in each packet, significant performance degradations reportedly occur in high traffic areas.

Encryption

Encryption is an ancient technique that was first used by the Romans. Simply stated, it is the scrambling of transmitted text using a set of rules (algorithms) known to the recipient, but hopefully to no one else. The recipient uses the same set of rules in reverse to unscramble the coded text and read the intended message.

There are two classes of encryption algorithm. They are:

Symmetric key. A symmetric key algorithm is one where the same key is used both to encode and decode the message. The most popular symmetric key algorithm is the Data Encryption Standard (DES), whose major advantage is that

it is fast and widely implemented. Its major limitations are its relatively small key size (56 bits), which weakens the security of the algorithm, and the need to exchange a secret key between the communicating parties before secure communication can be established. A variant of DES, Triple-DES or 3DES is based on using DES three times (normally in an encrypt-decrypt-encrypt sequence with three different, unrelated keys). Many people consider Triple-DES to be much safer than DES.

Asymmetric key. Some algorithms require additional codes to complete their calculations. These algorithms are difficult to reverse even when the algorithm, the key, and the encrypted data are all available, unless some other piece of information is known. A public key encryption system is based on an algorithm of this type. Two keys are generated. One is kept private and the other can be made public. Two systems that want to hold a secure conversation can exchange their public keys. When one system sends to the other, it will encrypt the message using the other system's public key. Even though an attacker might observe the exchange of keys and an encrypted message, the irreversibility of the public key algorithm ensures that the data is secure.

The most popular public key algorithm is the RSA algorithm. (The name RSA is derived from the first letters of the surnames of the algorithms inventors, Ron Rivest, Adi Shamir, and Leonard Adleman).

Although public key algorithms solve the key distribution problem, they are much slower than symmetric key algorithms. When implemented in hardware, DES is up to 10,000 times faster than RSA. If efficiency is required, a public key system can be used to securely exchange symmetric keys for the bulk of the data transfer.

Firewalls

The term "firewall" has become a generic term that encompasses a spectrum of technologies intended to provide protection from network attacks. Screening routers, application gateways, proxy servers, authentication servers, are all examples of firewalls in use today. It is possible, and often desirable, to combine these different technologies according to the needs of the organization and their budget limitations.

Application Gateways

All packets are addressed to an application on the gateway that relays the packets between the two communication points. In most application gateway implementations, additional packet filter machines are required to control and screen the traffic between the gateway and the networks. Typically, this involves the use of bastion hosts. Bastion hosts are secure but inefficient, since they are not transparent to users and applications.

Screening Routers

One of the most cost efficient and ubiquitous tools for securing networks are screening routers, sometimes known as a packet filtering routers. They allow known users (known by their IP addresses) to connect to specified applications (determined by their port address), thereby limiting the connections of those users allowed to enter through the firewall, and completely denying any connections to those not authorized to access any applications.

Authenticating Servers

Authenticating Servers are often used in conjunction with Screening Routers to provide authentication services, thus verifying that those users who claim to be originating from valid addresses are, in fact, who they say they are.