Security Issues on the Internet - A Technical Overview

M.Montebello, D.Attard, G.Attard

Computer Science and Artificial
Intelligence Department,
University of Malta,
Malta

Abstract

Increasingly, organizations are connecting to the Internet to establish a business and electronic commerce presence and to access information rapidly. When an organization's network is connected to the Internet without adequate security measures in place, it becomes vulnerable to attacks from external adversaries, unable to prevent many forms of undesirable access to its network, systems, and information assets. The risks include loss of confidentiality of business information, loss of availability of mission-critical services, exposure of critical data, legal liability, and vandalism of public information services.

Available technologies provide effective tools to manage the organizational networks' risk by providing access control mechanisms that can implement complex security policies. Some of the most popular and widely employed techniques are firewalls, virtual private networks and intrusion detection managers.

This paper will be looking at the technicalities involved in these techniques together with an analytic critique of their performance and of potential pitfalls.

1. Introduction

With the onward rush of electronic commerce on the Internet, there is widespread concern, particularly in the commercial sector, about security on the Web. Customers want to be reassured about the safety of entrusting their credit card numbers to a Web form. Companies would like to know that they can rely on the Web as a secure medium for business transactions. Software developers have taken note of the needs, and Microsoft and Netscape have incorporated cryptography software into their respective Web browsers to facilitate secure transactions and messaging.

The Internet may feel like a place where online users can roam anonymously and privately, especially if they access the WWW from their own home. However, with every foray onto the Internet, their computer and other computers actively exchange information. So just how private and secure are these communications? That depends primarily on the sites they visit and their Web browser's security features. Security on the Internet is jeopardized when user either browse, shop or carry out their banking needs.

1.1 Browsing

Browsing is the basic task that online users on the Internet perform using a Web browser. A Web Browser is an application software which allows users to browse Web pages that are hosted over a whole network of computer servers connected together to for the Internet. The shared information resident on these servers is what is known as the World Wide Web of knowledge (WWW). During browsing sessions users' computers and their personal information can be breached in various subtle ways as they are connected to a network with other people worldwide - the Internet. Depending on the World Wide Web sites they visit and the tasks they perform on the Internet, they may encounter and run, sometimes even unknowingly, a virus or other program that can harm their system or release private information to others. Some of these are expanded below:

1.1.1 Viruses: In order to protect themselves from viruses on the Internet, users should not download files from sources that they don't know are safe. Viruses usually are hidden in programs and activated when the programs run. They also can be attached to certain other types of executable files such as special-action Web files and video files. Generally, when download of a type of file that could contain a virus is going to occur, the browser will display a warning and ask whether the user wants to open the file or save it to disk. If confident that the file comes from a trustworthy source, then it may be opened. If unsure, then scan file with an anti-virus or cancel the download. There are many anti-virus programs that can scan a computer for viruses, inoculate against known viruses, and maybe even repair damage caused by a virus. To get the benefit of such a program, make sure it runs

as recommended. It may be able to schedule it to automatically scan the computer on a regular basis and get updates to the program, which include information about new viruses.

1.1.2 Cookies: A cookie is a small amount of information stored on a users' computers by a Web site, information that the users' Web browser sends back to the site whenever they visit it again. Usually the cookie is designed to remind the site of information about the users themselves, such as their password for the site or the customized background color they chose so that their browsing is simplified. Cookies are common and usually harmless. They can not be used to take information about users or their computer that they have not provided. But they can be used by certain services to create a profile of users interests based on the sites they visit. Then advertisements on participating sites can be customized for them. The browser can be set to alert users that a server is attempting to set a cookie on the host computer, while all the cookies can be prohibited up front.

1.1.3 Privacy: Any Web site users visit can tell exactly who and where their Internet service provider is, what site was visited last, what Web browser software is be used, and what actions are performed while on a particular site. By asking users to register, a site can collect additional information from them, such as their name, e-mail address, postal address, income level, and interests. It's up to them whether to provide this. Theoretically, the postmasters and system administrators who relay electronic-mail messages could read users e-mail if they wanted to. But so many e-mail messages are sent each day that it is unlikely any particular message would be read. Still, users should know that employers have the right to monitor e-mail sent using the organization's computers, law enforcement authorities can monitor e-mail under certain circumstances, and courts can require users to produce e-mail that relates to a court case. So it is a good idea not to say anything in e-mail that would not have been said in public.

1.1.4 Newsgroups: Messages posted to Usenet newsgroups are available to anyone on the Internet, and they are archived and can be searched, so these communications are not secure and private at all. Also, Spam e-mailers, those who send mass e-mail messages, sometimes pick up e-mail addresses from newsgroups.

1.1.5 ActiveX, Java, and certificates: Even if a user does not intentionally download software from a Web site, elements of a site may download, run on the computer, and pose a potential security risk such as by unleashing a virus onto the system. ActiveX® technologies allow software to be distributed over the Internet, usually graphic items such as scrolling marquees on Web sites, are similar to small programs that are usually digitally signed by its creator. Then a certifying authority can certify the signature. A certificate is an assurance that the control was safe when it was designed and that it has not been tampered with since. The Web browser can be set to enable, disable, or prompt the user to decide what to do with ActiveX controls depending on whether they are labeled safe. Java is a computer language. Java-based mini-applications, or applets, can be downloaded from Web sites and run by Web browser software. Generally, these applets are limited in what they can do.

1.2 Shopping

Shopping over the Internet is achieved by ordering items that have been browsed by users on an electronic commerce site (e-Commerce) set up by a business. The users will be asked to supply credit card details and this is what plenty of users are uncomfortable with. Such an action is probably more secure than users handing their credit card to a waiter in a restaurant or give out their account number over the telephone when ordering products. This is so because of the fact that some e-Commerce sites employ secure servers to encrypt, or encode, all the transaction information so that, if intercepted, it can not be read. Additionally, certificates enhances the users' trust as it provides an online document that certifies the business's identity to ensure users that the information is going where they intend it to go.

1.3 Banking

Banking and investing over the Internet is a commonly practiced procedure even though the security issues are very high. Online banks and investment services use encryption to protect the information in users' transactions. Before information leaves the Web site's server for users' computer, or vice versa, it is transformed or encoded into a cryptic way. After it reaches the appropriate destination, it is decoded. While the information travels over the Internet, where it may be vulnerable to being intercepted by someone with malicious intentions, it is essentially gibberish. There exist a number of different ways to encode information, as Cryptography is a massive research area in its own right. Just to give a slight idea of how difficult it is to decipher encrypted information, if a secure server and Web browser uses a 128-bit encryption, there are three hundred billion trillion of possible keys to unlock the code for each unique transmission, and only one of them works. Most of the widely used Web Browsers let you know when encryption is in use by displaying a padlock icon along the bottom of the browser window.

The paper is organized in this way. A general introduction to security issues from a user accessing the WWW point of view has been given in the introduction. The next three sections deal with Internet security issues from an organizational point of view and the vulnerability they are exposed to like attacks from external adversaries, unable to prevent many forms of undesirable access to their network, systems, and information assets.

The risks include:

- Loss of confidentiality of business information,
- Loss of availability of mission-critical services,
- Exposure of critical data,
- Legal liability, and
- Vandalism of public information services.

The performance of these technologies, which provide the effective mechanisms to manage an organization's network risks, together with their potential pitfalls, will also be discussed. Finally, our conclusions are presented together with some recommendation and suggestions regarding security issues on the Internet.

2. Firewalls

A firewall is a system or group of systems that enforces an access control policy between two networks. The firewall can be thought of as a pair of mechanisms: one that exists to block traffic, and the other that exists to permit traffic. Some firewalls place a greater emphasis on blocking traffic, while others emphasize permitting traffic. The firewall's configuration is the mechanism for enforcing policies, imposes its policy on everything behind it.

Firewalls are frequently used to prevent unauthorized Internet users (malicious hackers, worms, virii or otherwise) from accessing private networks connected to the Internet, especially private company or business intranets. All messages entering or leaving the intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.

There are several types of firewall techniques:

Packet filter: Looks at each packet entering or leaving the network and accepts or rejects it based on user-defined rules. Packet filtering is fairly effective and transparent to users, but it is difficult to configure. In addition, it is susceptible to IP Spoofing.

Application gateway: Applies security mechanisms to specific applications, such as FTP and Telnet servers. This is very effective, but can impose a performance degradation.

Circuit-level gateway: Applies security mechanisms when a TCP or UDP connection is established. Once (a safe) connection has been made, packets can flow between the hosts without further checking.

Proxy server: Intercepts all messages entering and leaving the network. The proxy server effectively hides the true network addresses.

2.1 Packet Filtering

A packet filtering firewall consists of a list of acceptance and denial rules. These rules explicitly define which packets will and will not be allowed through the network interface. The firewall rules use the IP packet header fields to decide whether to route a packet through to its destination, to silently throw the packet away, or to block the packet and return an error condition to the sending machine. These rules are based on the specific network interface card and host IP address, the source and destination IP addresses, the TCP and UDP service ports, TCP connection state flags, the ICMP message types, and whether the packet is incoming or outgoing.

The overall idea is that one needs to very carefully control what passes between the Internet and the machine that is connected directly to the Internet. On the external interface to the Internet, one will individually filter what's coming in from the outside and what's going out from the machine as exactly and as explicitly as possible.

2.2 Port-Level Filtering

This kind of firewall denies all incoming traffic by default. Connections to specific ports are then enabled explicitly.

Port numbers from 0 to 1023 are *reserved* ports. Reserved ports are bound to a particular service. Some server listening on that port, e.g. a Web server, routes an incoming connection to one of these ports to the appropriate service.

Each individual connection between a given client and server, possibly just one in a set of simultaneous connections to that server, is uniquely identified by the source address, port number, and transport protocol of the client in conjunction with the server's IP address, famous service port number, and transport protocol.

The firewall must not only allow access to the reserved ports that are offering services on, but must also allow access to the unprivileged ports as part of the ongoing connection between the client and server.

2.3 Application-Level Gateway

Application gateways provide an extra level of security, but generally at the loss of transparency to applications. Each application supported on a firewall requires a unique program to accept the client application's data and relay it to the server. The extra security comes from the data path. The firewall is actually acting as a server to the client, and a client to the destination server. The firewall verifies that the application data is of a format that is expected, and can filter out any known security holes. Extra authentication, logging of information, and even conversion can take place on the firewall.

Some advantages of application-level gateway:

- Fine-grained control of connections is possible, including filtering based on the user who originated the connection and the commands or operations that will be executed.
- Direct connections from the external network to the internal network are not permitted; all connections are passed to the proxies.
- Details of the internal network, including host names and IP addresses, can be hidden from the external network.
- Comprehensive logging and reporting systems can be implemented. Because more information is available about each connection, the firewall can write more detailed and more useful information to log files.
- Alarm system can be implemented, so that the firewall triggers real-time alarms in response to events that are regarded as potentially suspicious or hostile.

2.4 Circuit-Level Gateways

A circuit-level gateway relays TCP connections but does no extra processing or filtering of the protocol. For example, the TELNET application gateway example provided here would be an example of a circuit-level gateway, since once the connection between the source and destination is established, the firewall simply passes bytes between the systems.

2.5 Proxy server

Proxies are store-and-forward caches. When an application such as a web browser is configured to use a proxy, it never connects to the URL. Instead, it always connects to the proxy server, and asks it to get the URL for the application requesting it. Proxies isolate the user from connecting to the Internet. A proxy server receives a request for an Internet service (such as a Web page request) from a user. If it passes filtering requirements, the proxy server, assuming it is also a cache server, looks in its local cache of previously downloaded Web pages. If it finds the page, it returns it to the user without needing to forward the request to the Internet. If the page is not in the cache, the proxy server, acting as a client on behalf of the user, uses one of its own IP addresses to request the page from the server out on the Internet. When the page is returned, the proxy server relates it to the original request and forwards it on to the user. To the user, the proxy server is invisible; all Internet requests and returned responses appear to be directly with the addressed Internet server. (The proxy is not quite invisible; its IP address has to be specified as a configuration option to the browser or other protocol program.) An advantage of using a proxy server is that its cache can serve all users. If one or more Internet sites are frequently requested, these are likely to be in the proxy's cache, which will improve user response time. In fact, there are special servers called cache servers.

The functions of proxy, firewall, and caching can be in separate server programs or combined in a single package. Different server programs can be in different computers. For example, a proxy server may in the same machine with a firewall server or it may be on a separate server and forward requests through the firewall. There are different types of proxy servers with different features, some are anonymous proxies, that are used to hide the real IP address and some are used to filter sites, which contain material that may be unsuitable for people to view. When connecting to a web site, the true IP address will not be shown, but the proxy servers IP will, this does not mean that you are completely anonymous. The proxy server will have logs of IP's that used the proxy server and

the times. A proxy server can be used if you have some users who you want to restrict the sites they are viewing. It can be used to hide the user's IP which is useful because it means hackers can not get info about the user when using it. They will only get the proxy servers IP. Proxy servers are not hard to set up, no hardware or software is needed, applications that use the Internet must connect through it.

In practice, many firewalls use two or more of these techniques at once. What the firewall does depends on the way it is configured. Some firewalls permit only email traffic through them, thereby protecting the network against any attacks other than attacks against the email service. Other firewalls provide less strict protections, and block services that are known to be problems. Generally, firewalls are configured to protect against unauthenticated interactive logins from the "outside" (untrusted) world. This, more than anything, helps prevent vandals from logging into machines onto a network. More elaborate firewalls block traffic from the outside to the inside, but permit users on the inside to communicate freely with the outside. Firewalls are also important since they can provide a single point where security and audit can be imposed. The firewall can act as an effective tracing tool. Firewalls provide an important logging and auditing function. Often they provide summaries to the administrator about what kinds and amount of traffic passed through it, how many attempts there were to break into it, etc. Sengstack specifically recognizes the benefits of employing a firewall and states that the perfect personal firewall would be inexpensive and easy to install and use, would offer clearly explained configuration options, would hide all ports to make a PC invisible to scans, would protect a system from all attacks, would track all potential and actual threats, would immediately alert a user of serious attacks, and would ensure nothing unauthorized entered or left the user's PC.

2.6 Network firewall

A network firewall is one that is used to prevent access to a private network. It can also be used to limit traffic going out of the private network. The firewall has policies for accepting or rejecting connections, depending on a number of criteria specified by the user that handles the installation and configuration of the firewall. These kinds of firewalls are essential to many companies and businesses, because they protect the private networks from any intrusions and thus prevent and loss of important information.

2.7 Personal Firewall

Private users who are running computers that are permanently connected to the Internet e.g. using cable or ADSL usually use a personal firewall. Since these computers usually have a static IP, these become an easy target for hackers, Trojan horses, and other security attacks. Since the IP does not change frequently, if the computer is not behind a personal firewall it becomes a sitting duck. The personal firewall software provides the mechanism to prevent unauthorized to a user's computer. Attacks are usually done using scripts written purposely to scan certain ports that are known to have security flaws, so if the computer is found to be protected, the hacker usually leaves that computer alone.

3. Virtual Private Networks

A virtual private network (VPN) allows a company to securely extend its private intranet over the existing framework of a public network such as the Internet. With VPN, a company can control network traffic while providing important security features such as authentication and data privacy. VPN support uses the IP Security Architecture (IPSec) open framework. IPSec is unique in that it provides base security functions for the Internet, as well as furnishes flexible building blocks from which robust, secure virtual private networks can be constructed. VPN also supports what are called virtual lines; they provide cost-effective access for remote users by allowing the home network server to manage the IP address assigned to the remote user. In addition, these connections provide secure access to an organization's system or network when used in conjunction with IPSec. The basic idea of VPN-based extranets is to use the access control and authentication services with a VPN implementation to deny or grant customers, trading partners and business associates access to specific information that they may need to conduct business. With a VPN-based extranet application, the outside party would get to the corporate firewall by tunneling across the Internet or a service provider's network. The ability to get behind the firewall is controlled by the VPN access control services.

It's difficult to estimate the cost savings of using a VPN vs. another networking technology for extranets. For many companies, VPN-based extranets simply allow them to do business they could not do before like VPN authentication and access control services which are used to manage different levels of access, encryption, authentication and access control services to segment populations on a corporate network or intranet.

In many situations, companies need to ensure the confidentiality of data. For instance, a human resources department might want to let employees check on vacation time, but not be able to see performance reviews. Or a national sales manager might be granted access to the sales performance records of all sales associates, while each associate only has access to his or her own records. VPNs can help an IT manager establish and manage these levels of access.

Using VPN technology to control access to data for different groups of workers solves some problems that IT managers have faced for a long time.

For example, many IT managers have tried to segment user populations using virtual LAN (VLAN) technology. With VLANs, the idea is that a manager can quickly create ad hoc groups of workers who appear to be on a single LAN segment. A manager can dynamically assign users to specific groups and restrict others from any one group. The problem that a large number of IT managers encounter with VLANs is that many approaches are proprietary and therefore do not work in mixed environments where hubs and switches from multiple vendors are used. VPNs can cut across the mixed-equipment environment by using IP-based tunnels between a user's workstation and a server. The traffic between the two devices would be encrypted, which helps ensure confidentiality. So VPNs create an environment that is analogous to physically segmenting users on distinct LAN segments. By far the most exciting thing about VPNs is that all four applications are not mutually exclusive. A company could deploy a VPN to link its branch offices, then expand the access to single remote users and ultimately open up the network to outsiders, all using the same equipment and services. Once the connectivity needs of remote and outside users are satisfied, the installed equipment can also be used to segment user groups on the corporate network.

4. Intrusion Detection Managers

Computer intrusion detection managers are primarily designed to protect the availability, confidentiality and integrity of critical information infrastructures. These operations protect information infrastructures against denial of service attacks, unauthorized disclosure of information, and the modification or destruction of data. The automated detection and immediate reporting of these events are required to respond to information attacks against networks and computers. In a nutshell, the basic approaches to intrusion detection today may be summarized as known pattern templates, threatening behavior templates, traffic analysis, statistical-anomaly detection *and* state-based detection.

Computer intrusion detection systems were introduced in the mid-1980's to compliment conventional approaches to computer security. The most popular intrusion detection manager is D.Denning's seminal intrusion detection model that is built on host-based subject profiles, systems objects, audit logs, anomaly records and activity rules. The underlying intrusion detection model is a rules-based pattern matching system where audits are matched against subject profiles to detect computer misuse based on logins, program executions, and file access. There are other intrusion detection systems based on the Denning model and an excellent survey of these systems has been done by Mukherjee et al. whereby the basic detection algorithms used in these systems is highlighted, namely, weighted functions to detect deviations from normal usage, covariance-matrix based approaches for normal usage profiling, rules-based expert systems approach to detect security events.

The second leading technical approach to present-day intrusion detection is multi-host network-based. Heberlein *et al.* extended the Denning model to traffic-analysis on ethernet based networks with the network security monitor framework. This was further extended with the distributed intrusion detection system that combined host-based intrusion detection with network traffic monitoring. Current commercial intrusion detection system such as real secure and computer misuse detection system have distributed architectures using either rule-based detection, statistical-anomaly detection, or both.

A significant challenge remains for intrusion detection system designers to combine data and information from numerous heterogeneous distributed agents (and managers) into a coherent process that can be used to evaluate the security of cyberspace.

Conclusion

Many classic security problems, such as perimeter and host security, have become well defined and are routinely addressed by a wide range of product offerings, however computer and network attacks are still on the rise. Effectively combating these attacks is a network and security management discipline with emerging strategies and solutions. This paper underlined three security techniques, namely, firewalls, virtual private networks and intrusion detection systems, as best applied precautions taken against such threats.

Firewalls are specially designed devices that control the spreading of a network threat especially over the Internet. They intentionally impede an untrusted person from doing damage to an organization's networks. It is a method or device that regulates the level of trust between two or more networks. A firewall can consist of software, hardware or a combination of both.

Virtual private networks have become an accepted as a way to increase security in an insecure world, and to leverage the cost-effectiveness of the public Internet as a site-to-site, Extranet, and remote access network. The design and deployment of these networks is becoming an emerging industry practice.

The current state-of-the-art of intrusion detection systems is relatively primitive with respect to the recent explosion in computer communications, cyberspace, and electronic commerce. Organizations fully realize that cyberspace is a complex realm of vital information flows with both enabling and inhibiting technical factors. Identifying, tracking, classifying, and assessing hostile and inhibiting activities in this ever growing complex dimension is an enormous and fascinating technical challenge.

References

J.Sengstack, *Make Your PC Hacker Proof*, PC World, July, 2000

D.Denning, *An Intrusion-Detection Model*, IEEE Transactions on Software Engineering, Vol. SE-13, No. 2, pp. 222-232, February 1987.

Mukherjee, Heberlein, L., and Levitt, K., *Network Intrusion Detection*, IEEE Network Magazine, Vol. 8. No. 3, pp. 26-41, May/June 1994.

S.Snapp., *A System for Distributed Intrusion Detection*, Proceedings of IEEE COMPCON, pp. 170-176, March 1991.