COMPUTER SECURITY

Introduction

The drastic growth in the popularity of the Internet and in the use of communication between computers makes the subject of computer security very relevant. The reasons behind hacking keep changing. If once hackers had broken into computers in order to satisfy their curiosity and pump up their ego, nowadays more and more hackings are intended to steal information or damage systems for purposes of profit or revenge. Many times our basic lack of knowledge of the dangers that lie in the network, such as weak passwords or trusting any e-mail that reaches our mailbox, is what makes it easier For hackers. Nowadays, the web is crawling with viruses, worms, Trojan horses and the like (all the above mentioned concepts are explained below in our site) so that if your computer is not well secured and you are falling behind on your updates, catching one of them is almost certain.

Overview

It is hard to believe, but the first hacking attack happened 40 years ago. Since then the "Art" of computer intrusion has been involving and became much more sophisticated.

History of Intrusions

Prehistory (before 1969)

1945: Discovery of moth, captured between cables in a Navy computer, leads to birth of a two new terms; "Bug", a term used since the late 19th century to refer to problems with electrical devices; "Debugging", a term used to describe efforts to fix computer problems.

1949: Hungarian scientist John von Neumann invents the theory of self-replicating programs. His work served as foundation for computers' "memory".

1963: ASCII (= "American Standard Code for Information Interchange") code was Born. ASCII is a primitive computer language that allows machines produced by different manufacturers to exchange data.

1964: American government announced war against "phone freaks," which use "blue boxes" as tone generators to make free phone calls. To discover the identity of the group, 33 million telephone calls had been monitored. AT&T performed 200 convictions by the time the investigation ended in 1970.

1969: UNIX - the first multi-tasking operating system was developed. The ARPA company developed first network, used by government research groups and universities - ARPANET. ARPANET is the ancestor of the Internet.

Elder Days (1970-1979)

1972: John Draper using simplest whistlest unlocked AT&T's phone network, and allowed free calls and manipulation of the network.

1974: Telenet, a commercial version of ARPANET.

1979: First computer "worm". Term "worm" referred to short program that scanned a network for idle processors. The purpose was to provide more efficient computer use. That worm is the ancestor of modern worm - destructive computer viruses that modify or erase data on computers.

The Golden Age (1980-1987)

1981: > First "personal computer" was created by IBM. It was stand-alone computer, with a CPU, software, memory, utilities and storage.

1983: FBI arrested group of young hackers who break into several U.S. government networks using only an personal computer and a modem. "Computer virus" was at first used to describe a computer program that can "affect other computer programs by modifying them in such a way as to include a (possibly evolved) copy of itself."

1986: One of the first viruses,"The Brain", was released by programmers in Pakistan.

Crackdown (1988-1993):

1988: 23 years old Robert Morris creates a worm that penetrates ARPANET computers. By self copying, its program succeeds to flood the memory of attacked computers and disable more than 6,000 machines on the network. He was caught and sentenced to three years probation.

1991: Symantec releases the Norton Anti-Virus software.

Zero Tolerance (1994-1998): Public changed his attitude to hackers. Hackers stopped being only of government concern, now also the private user was terrified of attack actions on his personal computer.

1994: "Good Times" virus was spread by e-mail, by taking advantage of users' naivety. The virus was spread by e-mail warning its recipients from any messages with the phrase "Good Times" in the subject line. In addition the e-mail included recommendation to forward it ahead to warn as many people as possible. Of course the letter itself contained the virus, which erased a recipient's hard drive.

1995: Microsoft releases Windows 95. Anti-virus companies worried that the operating system would be resistant to viruses. Already in the same year new viruses appeared demonstrating that the new Windows operating system can be intruded.

1998: More than 500 military, government and private sector computer with Sun Solaris operating system, were penetrated by two teenagers. It was the first time that the government really realizes what are the how great the danger is.

1999: Thousands of computers were infected by "Melissa" virus, which spread with alarming speed. "Melissa" caused damage estimated in $80 million. The virus spread by sending copies of itself to the first 50 names listed in the recipient's Outlook e-mail address book.

Hack 2K (1999+)

2000: Millions of computers are getting infected by the "I Love You" virus. The virus sends passwrds and usernames that are stored on infected computers to the virus's creator. The creator of that virus was a young Filipino student. The creator was never arrested since no law against spreading computer viruses was legislated in the Philippines.

2000: Many Web sites, including Yahoo, eBay, Amazon and others high-profile Web sites, are knocked offline for several hours. The cause - a series of "distributed denial of service attacks" (DDOS attacks), which is a flood of traffic generated from hundreds of computers simultaneously. The source of the attack is a computer at the University of California - Santa Barbara.

2001: The "Anna Kournikova" virus is spreading by sending itself to every person listed in a victim's address book. The mail itself includes a digital picture of the young tennis player.

2001: The Code Red worm is a self-replicating program that infects many systems running Microsoft Windows NT and Windows 200 server. The worm sends e-mails from the infected computer to an IP address designated in the program. The target is to knock off the targeted system. One of the targets of the worm was the White House.

2001: The "Nimda" virus is spreading and infects hundreds of thousands of computers around thew world, several days after the Sept. 11 attacks. The virus is considered one of the most sophisticated.

2002: The largest ever DoS attack hits all "root" servers of the core Internet infrastructure. All of the 13 servers are being attacked. No slowdowns or outages are experienced due to safeguards built into the Internet's architecture.

2003: Hindreds of thousands of computers are being infected in less than three hours by the "Slammer" worm (known as the fastest-spreading computer worm). The worm caused havoc on businesses worldwide, knocking cash machines offline and delaying airline flights.

2004: The MyDoom worm uses "social engineering techniques" to persuade people to open an e-mail attackment that contains the worm. It becomes the fastest-spreading e-mail worm, but cause very little damage.

Threats to Computer System

Overview

The base principle of computer's attacks is exploit system vulnerability. Each system has its weak points and the hackers just know how take an advantage of them. They try best-known flaws with the suitable attack tools and in the most cases to there patient there is a reward.

In this chapter we present a broad view of different kinds of computer threats, selected based on their significance in our world today

Stealing Passwords

To enter a computer system, you probably should know the login command. Nearly all computers ask from the user to supply correct password within limited number of tries.

Login programs have been developed throw the years. For example in early years systems were just storing the login password in a file as is. Next stage was to save the password in confidential file, that way to keep the access only to users who knew its name. This method wasn't much better than the previous one, because also the security system's directory command would not list that filename, a system call does.

A lot of computers penetrations were caused due to entire password system weakness. The most common problem is weak passwords. People just tend to choose bad passwords. There are many researches on the subjects, and they all have the same conclusion: "There is a good chance to succeed in password guessing". Those researches do not say that all the picked passwords are poor, but one week password is more than enough for the intruder.

Password-guessing attacks take three basic forms:

- Log in using guessing username and password. The attacker may fail on his first tries, but all to often, one will work - and that is enough to break your computer defense. Only few operating systems will try to resist the attack from the inside. To solve the problem all, system should limit the user and allow only very small number of login attempts. The solution is to log the failures in some file and notify the account's owner of failed login attempts. As simply as its sounds, there are not many systems with good logging mechanism.
- Reusing already stolen password files to brake into new system. This kind of attack is usually very successful, because user seldom tends to reuse passwords.
- "Listen" to a legitimate terminal session and write down the used password. In this kind of attack, the quality of your password cannot rescue you.

There are number of self evident conclusions from all that was said above, but the most important one is that choosing a password is something that you should pick with serious consideration.

Social Engineering

Social Engineering = "technique of circumventing technological security measures by manipulating people to disclose crucial authentication information" (definition).

For example, event that took place at Bezeq:

"Hi, this is David Shelly, the sys admin. Someone called me about a problem with the "Is" command. I need you to change the password on my login on your machine, so I could fix the problem; it's been a while since I've used it."

"No problem."

So simple, with one call and some self-confidence, the attackers got the password.

Another example of Social Engineering is mail-spoofing. Message is sent (apparently from a system administrator) to some company worker asking to run some "test program" that prompts for a password or messages to system administrator (apparently from some inner worker).

The only way to overcome these kinds of threats is to perform strong authentication, for example "three-way handshake" (message-reply-ack of receiving the reply) is a good way to verify who the sender of the message is.

Many times people that are not aware to all threats concealed in the Net,are responsible for reproducing social engineering attacks. I am sure you have received message from your friend (apparently), with a worn from some file, for example system.exe, to be a virus. The e-mail probably said that you should erase that file without delay. It's a hoax. If you follow its recommendation you may damage your computer. There were a lot of such incidents, people tend to trust their friends.

Authentication Failures

Failure in authentication mechanism means that in some point on the authentication process, the system makes a mistake, and the defense wall is broken.

For example , your bank web page, in which your are being ask to identify yourself using username and password, could be duplicated by the attacker.So you will enter your password ,like you always do, but this time directly to the attacker's hands.

Another example is source-address validation method of authentication (only specific host are allowed to communicate with the system). Hackers just use "rpcbind" command to retransmit certain requests, causing the server to believe that legal user is trying to access the system. In this example even use of cryptographic authentication won't help.

The last example is attacker method called "Authentication Races": Let's assume that we know the length of the password and it contains only digits (0-9). The attacker creates ten connections to the desired service and waits. When some valid user will start typing his password, the attacker (attack program) will pass the valid digits throw all ten connections, creating situation in which ten host are trying to connect the system simultaneously. Before the valid user will manage to enter the last digit, the attacker will send different 10 digits (0-9) (different digits per each connection). Of course the computer is much faster than the human been, and one of the connection will be validated. In many cases the valid user will be rejected, because systems tend to allow only single login with each password. *The same trick could be done with characters instead of letters just with more established connections.

Viruses

Virus is a small program that copies itself to other programs or files. For example, a virus might attach itself to an executable program, such as a spreadsheet. Each time the spreadsheet program is executed, the virus is also executed. Computer viruses are so called because they share some of the characteristics of biological viruses. A computer virus passes from computer to computer like a biological virus passes from person to person. As a biological virus uses existing cells to reproduce itself, a computer virus needs an executable program to replicate itself.

Viruses may not immediately go on a rampage as soon as they infect a system. A virus may lie dormant for some time, and then get triggered on some specific date. But even while the virus is dormant, the machine is still infected.

The effect of a virus can vary from simply displaying irritating messages (for example, the WM97/Class-D virus, which repeatedly displays messages such as "I think 'username' is a big stupid jerk") to deleting specific files or formatting your hard drive (for example, the CIH virus, which tries to overwrite the Flash BIOS, can cause permanent damage).

Traditionally, viruses spread through detachable media, like floppy disks. However, now that we live in a networked environment, viruses don't solely rely on floppies or any other detachable media to spread themselves. Viruses may easily spread through e-mail or other documents that you download from the Internet.

Worms

Worms are programs that replicate themselves from computer to computer but do not infect other programs. The basic damage caused by worms is in terms of lost CPU time and the hours spent trying to eliminate them from the system.

Unlike viruses, worms do not use a host file to spread themselves. But worms generally exist inside documents only. The difference between viruses and worms is in the manner in which both use the host file. Worms usually release the document to which they are attached. This document then travels from one computer to another across networks. This implies that if a document has a worm attached, then that document, as a whole, should be considered a worm. However, a virus spreads from one document to other.

An example of a worm is the code red worm, which was one of the most destructive worms that affected Microsoft Internet Information Servers (IIS) across the globe. It defaced Web sites with a message, "Hacked By Chinese," and it spread at an alarming rate. Some 90,000 computers running IIS on the Internet were affected in only about 4 hours. The effect was so intense that many organizations involved in Internet security, such as SANS, CERT, Microsoft, and even the FBI, issued alert bulletins to people all over the world.

Trojans

A Trojan is a destructive program that masquerades as a useful or benign program. The main aim of a Trojan is to gather information from the computer. An example of a Trojan can be a program that masquerades as a signon screen and grabs the user's password when it is typed. However, unlike viruses, Trojans don't replicate themselves.

Trojans generally appear as games or valid programs; therefore, users install Trojans thinking they are legitimate programs. For example, a user might receive an e-mail message from an unknown sender. The message has a game attached. When the user executes the game, the Trojan is executed. Trojans can be hidden in any software.

Some Trojans are "proof-of-concept" programs written to illustrate a particular system weakness without doing much else. Others might delete files or even allow attackers to gain control of the PC. The best-

known remote-access Trojans, such as Back Orifice 2000 and SubSeven, contain many "fun" features such as audio and video captures, keyboard and mouse control, and CD-ROM drive control. Other capabilities might include sniffing passwords, creating and killing processes and services and deleting files.

What is a DoS (Denial Of Service) attack?

A DoS attack can be defined as an attack designed to disrupt or completely deny legitimate users' access to networks, servers, services or other resource. DoS attacks come in two basic flavors: **target resource starvation** and **network bandwidth consumption**. As with everything, there are a couple of exceptions; but for the most part, a DoS attack generally falls into one of these two categories.

Many DoS attacks use specially crafted packets. These packets can often cause operating systems or applications to crash, stop responding to requests, or otherwise fall prey to the DoS attack. DoS attacks often employ packets that should not exist. Any normal user or service on a network would not send out these packets. Because operating system designers did not anticipate that the TCP/IP stack would ever see such packets, responses to them are unpredictable.

A *resource starvation attack* attempts to deny service to a particular machine or a particular service on a machine. An example of this would be an attack that keeps legitimate users from accessing the FTP service. When the attacker wants to shut down multiple targets or even the entire network, he can use a *bandwidth consumption attack*. That means that the attackers flood the network with so much bogus data that no legitimate data can be processed. To accomplish this the attacker must use something to generate a tremendous number of packets. Because it is inefficient to do this only from the attacker's workstation, attacks using amplification are used. Amplification attacks rely on one packet generating many responses.

Example of DoS attacks

Ping of Death

The Ping of Death is just a very large ICMP packet which cause fragmentation. When a host receives this packet it starts reassembling the packet. However, because the size of this packet is very big, it overflows the buffer that is reserved for packet reassembling. This might cause unpredictable results. Reboots and system hangs are examples of the results.

Building this kind of ICMP packet is very easy on most systems. One can use Windows NT to send such a packet by simply using this command line: ' ping -165527 -s 1 '. Many source code examples for any king of Unix platforms are widespread in the Internet. They present simple techniques to construct such large packets.

Today, this type of DoS attack is useless as most systems have patches that prevent big ICMP packets from working. Since this is a simple method to create a DoS attack, it was used alot, and that brought the whole DoS issue to the awareness of system administrators.

Syn Flood

Another classic attack is the Syn Flood attack. This attack exploits a flaw in the TCP/IP stack itself. This flaw is in the connection procedure - the three-way handshake. recall that the three-way handshake consists of a SYN packet from a source, followed by a SYNACK from the destination and completed when the source returning an ACK. But wat happens if the last stage doesn't occur? What if this ACK is never sent back ?.

This causes the target host to wait untill the last ACK is received (the host shouldn't reject the connection because the packet may arrive in the near future). But, when this occurs often enough, the targeted host might not be able to establish new connections since all of its resources are used.

Since the handshake is not completed, most of the time the source address in the packets is spoofed, Making it hard to trace the real source of this attack.

Destributed Denial-of-Service (DDoS)

DDoS attacks are DOS Attacks that come simultaneously from many hosts conscripted from all over the net. They work as follows:

First, the attacker uses common exploits to install a zombie program on as many machine as he can, all over the Internet, in many different administrative domains. The zombie binds to a port and waits for instructions. Second, The attacker installs a master program somewhere on the Internet. The master has a list of all of the locations of the zombies. The master then waits for instructions. After the attacker waits for the time to strike, the attacker sends a message to the master indicating the address of the target. The master then sends a message to each of the zombies with the address of the target. Then, the zombies flood the target with enough traffic to overwhelm it.

Defend against DoS attacks

It is not easy to defend against DoS attacks (and especially DDoS attacks), but there are few things that sysadmins can do: For defending against malformed packets type of attack (the Ping of Death style of attacks), the best line of defense is to keep the system patched up, and to put a firewall between the system and the Internet. That firewall should be patched up and recognize "strange looking" packets. To defend from bandwidth consumptions attacks, such as the SYN flood attack, There is not much a sys admin can do, but try to reduce trafic by filtering the bad packets. This can be done by characterizing the incoming packets or even to shutdown the attacking hosts. Neither of these, are easy to perfrom.

How does a hacker work

Hacking Basics

Most system intrusions can be divided to approximately four steps:

1. Learning the target environment. This step includes several passive and active techniques that allows the attacker to gather information on the goal network, hosts in that network, their Operating systems, The services they provide and so on. From the information that wasacquired, The attacker finds the weak spots in the system and start to develop a plan that exploits those weaknesses

2. Attacking the system. At this point, the attacker initiate a series of procedures that eventually allow him to access the system at anytime he wish. These procedures could be anything from FTP access to a send mail bug to logging in as a "regular" user. The goal of this, is to test the weaknesses, until he finds a crack in the defense that give him access to the system services.

3. Full system access. After the attacker found a crack, he can use an exploit that gives full access to the system. From this point, He can do anything. using the right tools, the attacker can get anything that is on that computer, Including password files, accounts, configuration files and so on.

4. In this fourth step, Backdoors are installed and traces are covered. The backdoor gives the attacker the means to enter the system again very easily and quickly. Some experienced hackers even patch the system to keep less experienced hackers out of the system. The other part of this step includes removing logs and trace that may indicate that unauthorized user entered the system.

Network Scanning

In order to attack a network, An attacker needs to learn about the network, its hosts and the network services. The most direct way is to scan the network and its hosts. An attacker can locate hosts directly, through network scanners, and indirectly, perhaps from DNS or inverse DNS information. They may find targets in the host files on other machines, from chat rooms, or even newspaper reports. Numerous programs are available for host and port scanning. The simplest ones are nearly trivial programs, easily written in few lines of Perl or C. ICMP pings are most common host detection probes. UDP packets are also a tool that is commonly used (usually use the *trace route* port range). An attacker may scan an entire net host by host, or they may send directed broadcast packets. Directed broadcast are more efficient, but are often blocked.

Once located, hosts may be fingerprinted to determine the operating system, version and even patch level. Programs can examine idiosyncrasies of the TCP/IP stack (and sometimes can crash some hosts). Hosts are also scanned for active ports. They seek active network services, and often identify the server software and versions. Port scanners can be very subtle. For example, if they send a TCP SYN packet, but follow the computer response with an RST to clear the connection insteadof sending an ACK to complete the three-way handshake, a normal kernel will not report the connection attempt to a user-level program. Carefully crafted TCP packets can also probe some firewalls without creating log entries.

Breaking In

There are three approaches to breaking into a host from the Internet: a. Exploit a security hole in the network services offered by the host b. Duplicate the credentials of an authorized user or c. Hijack an existing connection to the host.

In the early day of the Internet, the first two were most common; now we see all three. There are other ways to break into machines, such as social engeneering or gaining physical access to the console or host itself. Security flaws are numerous. They are announced by various CERT orgznizations and vendors, usually without details. The hacking community discovers their own security holes as well, and sometimes exchanges them.

Some techniques are well-known from the start, like the ability to sniff (using programs like Ethereal or others) Ethernets for passwords. Others have been found during code reviews. Passwords can be sniffed or guessed, and other authentication failures can be exploited to break into a host.

When found, exploits are often engineered for simplicity of use. So, a hacker can gather some tools in order to exercise security flaws, that will eventually give a way to breach to the system.

At this stage the attacker is almost in the system. All it has to do is to exploit the system vulnerabilities in order to break its defences. After the attacker is in the system he might install a backdoor to allow him to quickly access the host in the future.

Exploits and Buffer Overflows

A buffer is a contigous block of memory that is used to store large amounts of information within a process. Buffers are stored in the data section in the process and are constant in length during run-time.

A common bug in computer programs is caused when the process tries to store more information than a buffer was designed to hold, causing an overflow. Many programming languages such as C, do not perform bounds checking on variable accesses so these overflows doesn't create any warning messages. Instead, the excess data is written into the process' memory adjacent to the buffer. Sometimes this overflow only ruins other data variables of the process and cause data corruptions. But sometimes it writes over the process' stack that holds return addresses of the process functions. When the function that caused the

buffer overflow wants to return, it uses the return address that is stored in the stack. That return address is now invalid and that usually cause a segmentation fault as the process continue to run from a random location in the memory.

An attacker uses this effect in order to gain access for his function. A buffer overflow used by a hacker leads the process to continue running from an implanted code (instead a random location), allowing the hacker to gain the same privileges as the attacked process. When used on system processes, the buffer overflow allow the hacker system privileges. The code implanted by the hacker can now excute other processes (i.e, creating new users), change or view restricted files and more.

Backdoors

A backdoor can be anything that gives the hacker a way back into the system. It allows him to access the system while bypassing all the existing security and sometimes it even defeat any additional security patches that were added onto the system.

The way to create that simple access may vary: There are simple backdoors that include creating new privileged users accounts for the intruser needs. some more complex backdoors can bypass regular access completely. Some involve trojans such as a modified login program that allows the attacker a privileged access if he type in a special password.
Another way is to chain backdoors together so the hacker have multiple access methods.

The hacker may alter configuration files in order to create a system misconfiguration that allows greater access. By doing that he lowers the defences of the system in some areas.

Auditing, Accounting, and Logging

Auditing, accounting, logging are things used to create permanent or semi- permanent records of events on a system. these can record intrusion activities, sometimes in explicit and evidence-worthy detail. Therefore, potential intruders are aware of what record keeping is available (either as a regular feature of the system or as add-ons) and have possible methods for defeating such recordings.

logging includes writeing login and logout entries to simple text files. Failed logins are also written to these files. Those might record an attempt to break into the system.
Auditing and accounting record programs that were accessed, programs that were executed, attempted to be run and failed. Some also keep track of an individual's disk usage. When the administrator suspect that his system was attacked, these tools might help to reconstruct the attack.

After attacking the system, the attacker usually tries to delete any log files in the system that might reveal the attack. This might be hard since these logs are usually stored in directories generally protected from editing. Most of the time sys admins prefer not to rely on the default logging tools supplied by the operating system but install other tools that use their own directories and log format, making it harder for the attacker to erase his traces.