IBM® Systems Journal

© IBM Corporation

Vol 37, No. 1 - Internet Computing

[Table of contents: HTML, ASCII]

[This article: HTML, ASCII]

Feature article

0018-8670/98/$5.00 © 1998 IBM

Catapults and grappling hooks: The tools and techniques of information warfare

by A. Boulanger

Reprint Order No. G321-5666.

For years, "hackers" have broken into computer systems, and now an entire industry is dedicated to computer network security. Both hackers and computer security professionals have developed software tools for either breaking into systems or identifying potential security problems within computer networks. This software can be found on compromised systems as well as within the toolkits of legitimate "tiger" teams that operate with the consent of the network owners. This paper describes some of the current techniques and tools employed by the hacker underground in breaching the security of networked computers, focusing primarily on UNIX®-based hosts connected to TCP/IP networks.

As organizations become increasingly dependent on computer network technology, they also become increasingly vulnerable to losses, of both financial resources and reputation, resulting from security breaches within their computer and communications infrastructure.

Many of the federally funded organizations dedicated to computer security issues were formed in response to attacks on computer systems. One of the first groups, Carnegie Mellon University's Computer Emergency Response Team (CERT[1]), was formed following an incident in which thousands of computers connected to the Internet were broken into and many disabled. This was the result of a self-replicating computer program, developed by Robert T. Morris at Cornell University, commonly referred to as the "Internet Worm."[2,3]

Computer security has become a serious issue. The media have reported a substantial number of recent attacks on high profile sites, and the number of reported security-related incidents is on the rise. In 1996 the United States Department of Defense (DoD) reported an estimate of 250 000 attacks per year on its computer system and stated that the rate of attack is increasing by 100 percent annually.[4]

When conducting examinations of systems that have been successfully attacked, certain patterns emerge. Data recovered from both the attacked systems and the computers of the intruders reveal similarities in how the intruders target and attack their victims. It has become clear that many of the components of the attack are automated and facilitated through use of sophisticated software toolkits.

The data for this paper were collected through our penetration testing work at the IBM Global Security Analysis Lab (GSAL) and through the lab's involvement with other organizations on actual security-related incidents. Every system that we have examined contained either the attack tools themselves or evidence of their use. These toolkits have been developed by both legitimate security professionals and by members of the "hacker" community. Many of the tools and techniques developed for legitimate purposes are modified and misused by computer criminals to compromise security and obtain unauthorized access to networks around the world.

This paper surveys some of the tools and techniques that have been successfully employed by members of both the hacker and professional security communities to break into networked computer systems. It defines the categories of tools and techniques used and outlines attack scenarios in which these tools can be deployed.

Intrusion tools

The tools and techniques in this survey can be broken down into five distinct categories. Each category defines tools and techniques that have been developed to exploit a specific type of system vulnerability. These categories are scanners, remote exploits, local exploits, monitoring tools or sniffers, and stealth and backdoor tools.

Scanners. A scanner is a tool that has been developed to obtain information about a host or network. Historically, these tools consisted of a loose collection of scripts [5] developed by security-conscious system administrators or by system crackers to probe the networks and report security-related information. Scanners can be broken down into two basic categories: network auditing tools and host-based static auditing tools. Network auditing tools are used to scan a remote host or series of hosts on a network and report security-related vulnerabilities for each host. Host-based static auditing tools are used to scan a local host and report its security vulnerabilities.

Network auditing tools. In 1992, Christopher Klaus released the Internet Security Scanner (ISS), one of the first network auditing tools to include many of the common security tests in a single package. [6] In 1994 and 1995, Dan Farmer of Sun Microsystems, and Wietse Venema, a research scientist from the Eindhoven University of Technology, developed and released the Security Analysis Tool for Auditing Networks (SATAN).[7] SATAN was based on the ideas presented in an earlier paper by the same authors.[8] SATAN expanded the functionality of ISS by adding more tests. It was designed to be portable, allowing it to run on a larger variety of platforms. SATAN's popularity and ease of use resulted in a large number of unauthorized scans of computer systems by system crackers as well as by curious users. In response to SATAN's release, system administrators developed software that would detect the "signature" of a SATAN attack. Among the first scan detectors to be published was "Courtney," developed by Marvin Christensen of Lawrence Livermore Laboratories. [9]

Today system administrators, as well as system crackers, have many free and commercial network security auditing packages to choose from. All of these software packages have the same goal: to locate and report network security vulnerabilities. SATAN, for example, will scan a range of host network addresses and report the following information:

- Host machines on the network that respond (and can be communicated with)
- Servers available on the responding hosts
- Shared disks available through Network File System (NFS) support
- File access through Network Information Service (NIS, a distributed database for shared information)
- Remote execution capability
- Sendmail vulnerabilities (versions that can be tricked into running bad commands)
- Trivial File Transfer Protocol (TFTP) access and configuration (can be used to download password files)
- Remote shell access (provides ability to execute commands on another system without entering a password)
- Unrestricted X Window System** server (anyone can connect to the server and spy on its users--a good way to obtain passwords and "freak out" users, drawing roaches or "smiley faces" on their screens)
- Readable/writable File Transfer Protocol (FTP) directory (allows any user to upload commercial software or pornographic material to corporate computing systems)

If SATAN reports the existence of any of these exposures, then it is likely that the system, and subsequently the network, are vulnerable to an outside attack. These types of problems are very well known to the hacker community and the tools to exploit them are widely available on the World Wide Web (WWW), anonymous FTP sites, and underground bulletin boards. Many of the tools used to breach network security can be easily found by using any of the Web search engines.

Host-based static auditing tools. The other type of scanner used by the security community is the host-based static auditing tool. Originally developed to "harden" the security of a local system, these tools, or their components, have also been used by hackers to obtain unauthorized privileged access. In 1989, Dan Farmer released one of the first static auditing tools, the COPS package.[10] This tool consists of a collection of scripts that scan the local system seeking out and reporting security vulnerabilities. In 1992, Texas A&M developed and released the TIGER toolkit, which expanded upon the original ideas behind COPS and added greater functionality. [11] Both of these tools perform extensive system checks and report the following vulnerabilities:

- Permission problems in files, directories, and devices, allowing intruder access
- Poor, easy-to-guess passwords
- Poor security for password and group-definition files
- Known vulnerable services, for example anonymous FTP configuration, or improperly configured services
- Signs of past intrusions, particularly in key binary files

The static auditor is a valuable tool to both the hacker and the system administrator. If the hacker is able to get an unprivileged account on the system, the local scanner will point out common security weaknesses in the host that enable unauthorized privileged access.

Remote exploits. A remote exploit is a program, or method, that can be used, by a person who has no existing account, to penetrate a remote computer system. The vulnerability to remote exploits is the driving force behind the development and deployment of firewalls and network auditing tools. A firewall can be defined as a system or group of systems that enforces an access control policy between two networks.[12] The firewall protects the network by defining the external services it provides and hiding information about its internals. A firewall, when properly configured and maintained, helps to protect the internal network from remote attacks by minimizing its exposure to the outside world. Reducing the number of services available to the external network helps to reduce the vulnerability of the internal network to remote exploits.

Remote exploits are associated with services provided by computers in the network. Most services will open a communication channel and listen for incoming connection requests. In the case of *sendmail*, a program that processes electronic mail, the program will open a communications port and listen for incoming requests from other sendmail servers. A sendmail server accepts a connection and communicates with the client system on the network using SMTP (Simple Mail Transfer Protocol). [13] If the sendmail server has a security vulnerability that can be exploited through user-defined data, then the server's host is vulnerable to attack from unprivileged users on any connected system. This is the main reason that remote exploits are the most feared and dangerous, and therefore the most closely guarded, of all the tool sets.

A subcategory of the remote exploit is the protocol-based attack. A protocol attack is a tool used to gain unauthorized access to a computer system through the manipulation of the Transmission Control Protocol/Internet Protocol (TCP/IP) network protocol suite. Vulnerabilities in the TCP/IP protocol have been known for years. In 1985, R. T. Morris described a vulnerability in the TCP/IP protocol that would let a hostile system appear to have another host's address. [14] If the targeted system is using a protocol that relies on address-based authentication, a hostile host has the ability to subvert that authentication and access the target system as a trusted host. In 1989, Steve Bellovin published a paper that generalized this type of attack and reported other new security-related problems with TCP/IP protocol.[15] These vulnerabilities include session hijacking and IP spoofing (both UDP [User Datagram Protocol] and TCP), RIP- (Routing Information Protocol-) and ICMP- (Internet Common Message Protocol-) based attacks. The last reported high-profile incident using a protocol-based attack was the IP-spoofing attack allegedly launched by Kevin Mitnick in December of 1994.[16]

Local exploits. A local exploit is a tool or method used to gain unauthorized privileges on a computer system by a person who has an existing account. This existing account can either be legitimate, obtained through a remote exploit, or otherwise acquired (for example: traded with other hackers, captured from network traffic, by social engineering,[17] etc.). Most local exploits result from errors in a privileged program's software design or implementation that allow an unprivileged user to execute hostile commands at a privileged level or access and modify privileged data. [18] Once privileged access is attained, the hackers are in control of the system. On the majority of UNIX systems, the intruder is able to modify the system logs to hide illicit activities and install a "backdoor" that allows continued privileged, and unlogged, access to the system. On average, new local exploits are reported at over three times the rate of new remote vulnerabilities[19,20] and are widely available through security-related newsgroups, mailing lists, and Web sites. This is why it is considered to be a good practice for system administrators to utilize tools such as COPS and TIGER to help ensure that systems can withstand such attacks.

Monitoring tools. A monitoring tool allows a user to monitor the computer system and the network data. Intruders use this information to prepare attacks against other computer systems. This category of tools includes:

- Sniffers. A "sniffer" program monitors and logs network data. The network traffic that passes through a host's network interface usually contains user name-password pairs as well as other system information that would be useful to an intruder. Most systems do not encrypt the data that are transmitted on a computer network. A hacker with physical access to the network can plug in a sniffer, monitor the network traffic, and gain enough information to be able to access other systems on the network.
- Snoopers. A "snooper" program monitors a user's activities by snooping on terminal or terminal emulator sessions, monitoring process memory, or logging a user's keystrokes. By watching the victim's actions, an intruder can gain information, then use it to attack other systems on the network.

Stealth and backdoor tools. A stealth tool allows an unauthorized user to modify the system logs and eliminate all records relating to his or her activities. Stealth toolkits often include "backdoor" programs. These are modified, drop-in replacements of critical system binary code that provide authentication and system reporting services. Backdoor programs can:

- Provide continued, unlogged use of the system when activated (activation mechanism is often an encrypted password compiled into the program)
- Hide suspicious processes and files from the users and system administrators
- Report false system status to the users and system administrators
- Report false checksums for the modified programs

Deployment

With a collection of tools and "exploit scripts" in hand, the intruder can then move on to attack a computer network. Many intrusions are conducted against random targets where the main goal is to breach network security. These attacks, while common, are motivated by intellectual challenge rather than monetary gain. However, there is mounting evidence of a more focused type of attack on the networks of specific organizations for the purpose of fraud and espionage.

Penetration scenarios. Regardless of the intent or goal of the intruder, the attack will consist of a combination of one or more of the following penetration scenarios: blind remote attack, user-level attack, and physical attack.

The blind remote attack. The blind remote attack is a remote penetration attempt upon a computer or computer network where the intruder does not have valid account information or access, but knows the network address in either numeric or text form. This is the "classic" attack scenario: an unknown individual is illegally accessing a computer network. Most of the penetration testing conducted by security consultants includes, at the very least, a blind remote attack.

With only the address or name of the target system, the intruder will attempt to use network scanners and other methods to acquire security-related information about the system. After scanning and probing the system's defenses, the intruder will attempt to apply the right remote exploit from his or her toolkit to the vulnerable service. If the exploit is successful, the intruder will have at least user-level access to the computer system.

The user-level attack. A user-level attack is a penetration attempt into a computer system on which the intruder has user-level, or unprivileged, access. The exploited account can be legitimately acquired, as a customer or employee of the organization, or otherwise acquired by "sniffed" passwords, traded accounts, "shoulder surfing," blind remote attack, cracked passwords, social engineering, or default user accounts. The majority of financial losses resulting from breaches in network security are the result of "inside jobs," where a legitimate user attacks the network from the inside.

The first stage of this attack is to gain information about the computer system and its users. A local system scanner (COPS or TIGER) can be used to detect and report common security vulnerabilities. After scanning the local system, the intruder can apply the appropriate local exploit from his or her toolkit against that system. If successful, the intruder will have privileged access to the computer system and will then be able to compromise other systems on the network. If one site suffers a breach in security and its system is penetrated, there is a very good chance the intruder will gather enough information, through monitoring the system's data and network traffic, to gain unauthorized access into other machines on the network.

The physical attack. In the physical attack scenario an individual with physical access to the computer and the network equipment attempts to gain access into other networks and hosts. In this scenario, an intruder can plug a computer into the network and begin monitoring traffic. After collecting the logged network data, the intruder can use that information, either locally or remotely, to gain access into hosts on the network. Another common scenario is a physical attack on the host computer itself. With physical access to a system, it is very easy to gain entry. Many users leave their computers on, with their accounts logged in, when they leave the office. With physical access to these systems, the intruder can easily gain enough information to penetrate the network. If the targeted computer system has no active sessions, the intruder can then shut down and reboot the system. Under some system configurations, the intruder could then gain administrative privileges on the system, making it and other systems on the network vulnerable to attack.

In each scenario, the intruder performs steps in a sequence. These steps, or stages, form a "system penetration protocol." The seven stages of system penetration are:

1. Reconnaissance--gather information about the target system or network
2. Probe and attack--probe the system for weaknesses and deploy the tools
3. Toehold--exploit security weakness and gain entry into the system
4. Advancement--advance from an unprivileged account to a privileged account
5. Stealth--hide tracks; install a backdoor
6. Listening post--establish a listening post
7. Takeover--expand control from a single host to other hosts on network

Every intrusion will use some of the steps of this protocol. In the blind-remote scenario, the intruder would likely use, at the very least, the first three stages. The intruder would first attempt to gather information about the targeted system (stage 1). Then, using this information, the intruder would apply the remote exploit tools and techniques (stage 2) in an attempt to gain a toehold into the network. If the penetration attempt was successful and the toehold (stage 3) is that of a privileged account, the intruder can immediately begin covering his or her tracks and establishing a listening post (stages 5-7). If the toehold is that of an unprivileged account, the intruder would seek to obtain privileged access using a local exploit (stage 4). Once a privileged account is obtained, the intruder can proceed.

In the user-level attack scenario, the intruder has already achieved a toehold (stage 3) into the targeted network. This toehold could have been attained by methods ranging from user name and password guessing to cracking the password file obtained from the remote system. Cracking the password file is a very effective way to gain entry into a system. Once one password file has been obtained, the intruder is likely to guess approximately 25 percent of the remaining passwords.[21] In this scenario, the intruder's challenge is to obtain unauthorized privileges (stage 4). To accomplish this, the intruder obtains information about the local system (stage 1). Then, using local exploits, the intruder applies the appropriate tool (stage 2) and obtains unauthorized privileges (stage 4). Once the privileges are obtained, the intruder hides evidence using stealth toolkits and installs a series of backdoors to ensure future access. Now, the intruder can gain control of the network using the system monitoring tools. If the acquired account is privileged, as would be the case for a system administrator, the intruder has immediate access to all traffic and data on the system.

In the physical attack scenario, the intruder may follow the user-level attack scenario to find an active session on the system or reboot the system to gain administrator privileges. If the intruder is forced to monitor the physical network (stage 6), then the information gained (stage 1) can be applied to gain access into the system (stages 1-7). Having physical access to the computer system and its network hardware makes it very easy to compromise the system. It is for this reason that employers should provide physical protection to any sensitive computer systems.

The attack. The following example illustrates the application of tools and techniques and the penetration protocol on a targeted computer network. In this example, the intruder launches a blind remote attack on a computer network owned by the XYZ corporation, registered in the "xyzcorp.com" domain. The only information available at the start of this penetration is the name of the corporation.

Reconnaissance. The intruder wants to attack the computer systems owned by the XYZ corporation and begins the reconnaissance stage by searching the Internet for references to that corporation. If the XYZ corporation has an Internet connection, Web site, FTP site, or electronic mail service, then it is very likely that a Web browser will find a reference to the target's domain name. In our example, the search yields a domain name for the XYZ corporation, "xyzcorp. com." Using the domain name, the intruder can obtain more information through a variety of methods. One is the domain information groper, or "dig" program, a utility developed by Steve Hotz. [22] Our intruder uses it to attempt a "zone transfer" on the domain's name servers to get information about other machines within that domain.

We will assume that the attempt was successful, yielding a list of host names and network addresses from the targeted system. Now the intruder can gather information about the users on the system. Two excellent sources of information on the users within a particular network are the Web and newsgroups. Searches on network and domain names using the Web or searching the news hierarchy in a domain may yield a list of new hosts and a partial list of users on a system. The user list is very important; it may reveal user name-password combinations and possibly the domain's policy of determining user names. For example, if a search yields "From: bobr@host.xyzcorp. com (Bob Reilly)" from a news posting, the intruder can now attempt to open the account for that user name, guessing at the password. If the search also yields "From: sarag@hostb.xyzcorp.com (Sarah Gregory)," there is a chance that the user names for the entire system were constructed by concatenating a user's first name with the first letter of his or her surname. The intruder can attempt to guess additional user names and passwords, or search for the user name on chat channels (Internet Relay Chat [IRC], Web Chat) and attempt to acquire the user's personal information (name, address, phone number, etc.). With the personal information the intruder might then contact the user either by phone, electronic mail (e-mail), or chat, and acquire account information through persuasion (social engineering), or trick the user into running a hostile or "Trojan horse" program that could capture account information and send it back to the intruder.

In our example, at the end of the reconnaissance phase our intruder has the following information:

- Host name(s)
- Host address(es)
- Host owner
- Host machine type
- Host operating system
- Network owner
- Other hosts in the network
- Network configuration
- Other hosts trusted by the network
- Hosts outside the network
- List of users
- User-name assignment policy

Probe and attack. In this stage, the intruder begins probing the perimeter of the system's security for potential weaknesses. This is the most heavily automated stage of the penetration cycle. Toolkits recovered from compromised sites always include some type of scanner that enables the intruder to conduct security surveys on entire networks. Security professionals have taken scanning technology one step further to develop both public domain security scanners (SATAN) and commercial scanners (ISS). These automate the collection and reporting of security-related vulnerabilities of remote hosts and networks. This is, by far, the most dangerous phase for the intruder. The scans and probes are the activities most likely to be detected and logged by intrusion detection systems (if installed) that alert users and security-conscious system administrators.

Probing a system for security weaknesses is easily accomplished by determining what remote services each of the hosts is providing. Using a publicly available tool, "strobe," [23] an intruder can scan a host, or range of hosts, and generate a list of services offered by each host. In our example, "host.xyzcorp.com" is scanned using the strobe tool, and a list of services is generated.

The services of interest on this host are FTP, SMTP (for e-mail), finger, WWW, printer, and finally, xterm, the X Window System server. From his or her exploit toolkit, our intruder now looks for remote exploits against these services. First, the FTP server is checked for known vulnerabilities and configuration errors. Then the SMTP, or sendmail server, is probed and the service information containing the machine name and software version number noted. Sendmail, like most network services, often provides useful information to intruders, for example, the software name and version number, allowing the intruder to easily find the right exploit. If bogus information, or no information at all, is provided in the server's banner, the intruder's task is more complicated and the likelihood of detection is increased. Each of the services will be tested until a potential vulnerability is found. For our example we will assume that all the services tested are secure except the Web server. The WWW server on host.xyzcorp.com offers the vulnerable "phf" service, [24] and our intruder has a remote exploit against it.

The hostile command is executed on the server, yielding an X Window System terminal emulation on the intruder's display. A toehold into the targeted network has been obtained and the intruder advances to the next phase.

Toehold. A security weakness has been exploited and the system is compromised; the intruder has gained access into the system. If the user identification (UID) of the X Window System terminal is "root," the intruder moves on to the stealth phase. If the UID is for an unprivileged user, the intruder attempts to advance from the unprivileged account to a privileged, or administrative account.

Advancement. The intruder uses the information about the host, its operating system, and the services it provides to search the toolkit for the matching local exploit. In this example, the intruder has obtained a local display running a shell [25] on the remote server with UID "www." Now the intruder can use the local scanning tools (COPS/TIGER) to search and report configuration errors and other known vulnerabilities and apply local exploits from the toolkit. In this example, the local scan using COPS revealed the host to be an AIX* 3.2 (Advanced Interactive Executive) machine and vulnerable to the "tprof" exploit. The intruder can now advance from UID "www" to UID "root" and proceed to the next phase.

Now our intruder has successfully obtained the highest level of privileges and is in control of the targeted system. On most systems, any local file could be accessed and modified. Some malicious intruders look around for interesting data, then delete the entire file system. Most intruders retain their access to the compromised system, and move to the next stage: stealth.

Stealth. When the root is compromised, the intruder probably attempts to cover his or her tracks and avoid detection. The root UID has access to all of the files on the local system, so the intruder can now begin editing the log entries to remove evidence. In our example, the intruder checks the WWW server access logs for previous intrusions and then deletes all traces of the illicit activity. By replacing the system's binary code with modified versions that hide process and file information, as well as network connection information, the intruder removes all incriminating traces, and is ready to ensure future access to that system and to establish a listening post.

Listening post. The intruder now ensures continued, unlogged, and undetected access to the compromised system. Using one of the "root kit" packages from the toolkit, the intruder "patches" the programs on the system, to serve three main purposes. The first is to ensure that future activity will not be logged. The patched binaries contained in the root kit packages report false information on files, processes, and the status of the network interface to the administrators. The second purpose is to ensure continued, unlogged access to the system through a number of backdoors. The third is to establish a listening post for the network. The listening post is accomplished through a sniffer program that allows any privileged user to capture traffic on network interfaces that support "promiscuous mode."[26] If the targeted system does not have a promiscuous-mode-capable interface, the intruder is limited to monitoring the activity of each user on the local system. Network traffic contains sensitive e-mail and user name-password combinations for other systems and networks. By logging the interesting network traffic, the intruder can expand his or her area of control in the next phase.

Takeover. Using a combination of sniffed user name-password combinations and the toolkit of local and remote exploits, our intruder can attack other hosts on the network. Beginning from a single weakness in a single machine within a computer network, the intruder now expands the area of control. The installed backdoor ensures detection avoidance and continued, unlogged, privileged access to a series of hosts. The passwords that have been obtained from the listening posts provide all of the information required for obtaining future toeholds and root compromises. The intruder can use this information, compromise other machines, and rapidly advance through the network.

Conclusion

In 1995 the Computer Emergency Response Team (CERT) reported 1168 security-related incidents.[27] That year the United States Federal Bureau of Investigation (FBI) disclosed the results of their computer security survey, which showed that 40 percent of the surveyed sites experienced at least one unauthorized access.[28] CERT's 1996 figures show a significant increase in hacker activity, with 2573 security-related incidents.[29]

Organizations are now beginning to address seriously the issues of electronic and computer security. At one time most organizations would build a system and then, if it worked, install security precautions as an afterthought. Security concerns need to be addressed throughout the development and maintenance phases of each project. This need is evident in *Information Week*'s annual survey results published in October 1996. Of the organizations responding to the survey, 78 percent reported financial losses resulting from security breaches.[30] Many of these incidents could have been avoided, or at least minimized, if the operators of the attacked networks had taken some basic precautions. If an organization depends on its computer network for its daily operations, it should take the steps necessary to better secure its systems. If an organization has an Internet connection to its networks, installing a firewall can provide effective protection against most remote attacks. For internal system security, the organization can conduct random security audits of the internal network to lessen its exposure to local attacks. While no computer system is totally secure, applying some of these basic precautions can substantially reduce the possibility of a successful attack on an organization's vital assets.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of X Consortium, Inc. or X/Open Company, Ltd.

Cited references and notes

Accepted for publication September 30, 1997.

Reprint Order No. G321-5666.

[Journals home page | Subscribe/order | Current issue | Recent issues | Description]

IBM Home | Shop | Contact IBM | Search | Privacy | Legal