Global Information Assurance Certification Paper

Don't Let Hackers Gather Your Information

Robert Maheu

April 15, 2001

GSEC Practical Assignment Version 1.2b

Introduction:

Research is the starting point for almost every new project. By doing research you will gain valuable information about the subject in question. The same holds true in the world of hacking. A hacker will first attempt to gain information about a system before trying to break in. There are several avenues to take when it comes to gathering information.

Social engineering:
Having an individual perform an action or provide information just by asking them. This includes the acquisition of names, phone numbers, and possibly even passwords.

Dumpster diving:
Rummaging through the garbage to collect such things as manuals, software, and old computer printouts. Each of which can provide some form of information about a network or users on that network.

Network tools:
Software programs such as port scanners, sniffers, and network management. Although these tools can help administrators secure and manage their network, they can also be used with malicious intent.

This paper will focus on using network tools as a way of collecting information. It will also discuss basic ways in which to prevent this information from being obtained.

Port Scanning:

There are many applications that use the TCP/IP protocol. Most computers will use more than one of the applications in daily use. For example, a machine that is used for network management (SNMP) would more than likely be listening for SNMP traps, using Telnet and possibly even have a TFTP server running. There had to be a way for each datagram to get to the correct application. This is where the idea for ports comes into play.

According to the Internet Assigned Numbers Authority (IANA), the ports are divided into three categories (http://www.isi.edu/in-notes/iana/assignments/port-numbers):

a) The Well-Known Ports are those from 0 through 1023.
b) The Registered Ports are those from 1024 through 49151
c) The Dynamic and/or Private Ports are those from 49152 through 65535

Well-Known ports are assigned for specific use (such as FTP, Telnet, SMTP, etc.) and can not be used for other purposes. Registered and Dynamic ports can be used for any purpose (although Registered ports that are in use are assigned for one purpose). It's in this region that trojan horses such as Back Orifice and SubSeven live.

If one device wanted to talk to another using SNMP, the SNMP port has to be in an active or open state so that it can receive the datagram. Since this is the case, it is possible for other machines on a network to determine which ports are in an open state for a particular system. Port scanners are used for just this purpose. Because of the information that can be obtained from a simple port scan, it will probably be the first thing that a hacker will use on system.

The following is an NMAP session that was run on a test system. NMAP is a very popular port scanner with many features. It is currently a free tool available at www.insecure.org. Please note that the scans have been truncated to keep the list of open ports short.

Starting nmap V. 2.54BETA7 (www.insecure.org/nmap/)
Interesting ports on test.test.com (172.26.1.2):
(The 1504 ports scanned but not shown below are in state: closed)

Port	State	Service
7/tcp	open	echo
9/tcp	open	discard
13/tcp	open	daytime
19/tcp	open	chargen
21/tcp	open	ftp
23/tcp	open	telnet
25/tcp	open	smtp
37/tcp	open	time
79/tcp	open	finger
80/tcp	open	http
111/tcp	open	sunrpc

Starting nmap V. 2.54BETA7 (www.insecure.org/nmap/)
Interesting ports on test.test.com (172.26.1.2):
(The 1423 ports scanned but not shown below are in state: closed)

Port	State	Service
7/udp	open	echo
9/udp	open	discard
13/udp	open	daytime
19/udp	open	chargen
37/udp	open	time
42/udp	open	nameserver
69/udp	open	tftp

111/udp open sunrpc
161/udp open snmp
162/udp open snmptrap

Some of the interesting ports that are listed in the scan are SMTP(25), TFTP(69), and SNMP(161). It is possible to gain further information about a system or a network from just these ports. There are certainly other ports that can be used. However, these will be left for the reader to research on their own.

The following are some simple steps that can help reduce the amount of information that can be gathered by a port scan.

a) Run a port scan on your own system. You will be able to see if there are any ports open that you are not using. You should have as few open ports as possible. Take steps to close any ports that you do not need.

b) Use some form of a firewall. A firewall can be used to block specific services from going to/from parts of your network. For home use, there are several good personal firewalls. A good site for a quick review (and test) of personal firewalls is http://grc.com/lt/howtouse.htm

c) Install intrusion detection software. This will allow you to detect malicious activities on your network. Port scans can have specific signatures that can be detected.

SMTP – TCP Port 25:

It is possible to telnet to other ports. When you telnet to other ports there is the chance that you can get information about the application using it. The following example is a telnet session to port 25 (Simple Mail Transfer Protocol) of the same system that NMAP was run against. The first thing that you see when you telnet into the system is the actual mail software in use:

220 hurricane22 ESMTP Sendmail 8.11.0/8.11.0; Thu, 22 Mar 2001 09:17:15 –0500

Here it is seen that the software is Sendmail and that version 8.11.0 is being used. This information can be used to find security holes within the software. Typing 'help' at this point will provide the following list of commands:

214-2.0.0 This is sendmail version 8.11.0
214-2.0.0 Topics:
214-2.0.0 HELO EHLO MAIL RCPT DATA
214-2.0.0 RSET NOOP QUIT HELP VRFY
214-2.0.0 EXPN VERB ETRN DSN AUTH

```
214-2.0.0      STARTTLS
214-2.0.0 For more info use "HELP <topic>".
214-2.0.0 To report bugs in the implementation send email to
214-2.0.0      sendmail-bugs@sendmail.org.
214-2.0.0 For local information send email to Postmaster at your site.
214 2.0.0 End of HELP info
```

Using “HELP <topic>” on every command shows that two of them can be used to see if there is a user(s) configured for this system.

```
214-2.0.0 VRFY <recipient>
214-2.0.0      Verify an address.  If you want to see what it aliases
214-2.0.0      to, use EXPN instead.

214-2.0.0 EXPN <recipient>
214-2.0.0      Expand an address.  If the address indicates a mailing
214-2.0.0      list, return the contents of that list.
```

Running the VRFY command on the test system produced the following results:

```
vrfy root
250 2.1.5 root <root@hurricane22>
vrfy bmaheu
550 5.1.1 bmaheu... User unknown
vrfy maheu
250 2.1.5 Bob Maheu <maheu@hurricane22>
```

Verifying user names can be relatively easy. By doing a little research on a company you should be able to collect names of people that work there. More than likely a potential hacker will already have names from the company’s press releases and web site. Once they have this information they can start to guess the user names. There’s a pretty good chance that the user names will be in the form of first name, last initial or some variation of this scheme.

It is easy to stop these commands from confirming that a user exists. This version of Sendmail uses the sendmail.cf file for configuration. In this file you will need to look for the section titled ‘privacy flags’. You will then need to add ‘novrfy’ and ‘noexpn’ to the PrivacyOptions line.

```
# privacy flags
O PrivacyOptions=authwarnings,novrfy,noexpn
```

Once this is done, you will get the following output when trying to use VRFY and EXPN:

```
vrfy test
```

252 2.5.2 Cannot VRFY user; try RCPT to attempt delivery (or try finger)
expn test
502 5.7.0 Sorry, we do not allow this operation
Section 2.11 of RFC 2505, Anti-Spam Recommendations for SMTP MTAs, also suggests that these two commands be disabled to prevent spammers from verifying addresses. (http://www.ietf.org/rfc/rfc2505.txt?number=2505).

TFTP – UDP Port 69

Trivial File Transfer Protocol (TFTP) can be thought of as the little brother of FTP. Both are used to transfer files from one system to another. TFTP is used mostly to boot diskless workstations or to send new software images to network devices (such as routers). Some other differences are that TFTP uses less commands and has no user authentication. Since there is no user name or password required for TFTP, it may be possible to get any file off a system (as long as you have 'read' permissions on that file).

On the test station, that was incorrectly configured, it was possible to TFTP the /etc/passwd file to another station. This is the file that contains user names and passwords. Once this file is obtained, it is an easy task to crack the passwords with program likes Crack and John the Ripper (the only factor is time). As a side note, it took about 18 seconds on a 450Mhz PentiumII to crack the weak password 'startrek'. It took over 20 minutes to crack a stronger password of 'b08j8y5r' (the process was actually stopped before the password was cracked). When creating passwords, you should use a mix of alphanumeric characters along with special characters.

As one can imagine, TFTP can be a big security hole for a system. Luckily there are a few things that can be done to help reduce the risk.

a) TFTP can be configured to allow transfers to/from a specific location. On a Unix station this is usually done in the /etc/inetd.conf file. Search for the tftp line. It should look similar to the following:

   #tftp dgram udp6 wait root /usr/sbin/in.tftpd in.tftpd -s /tftpboot

   The # sign is a comment (which is a good thing since by default TFTP is disabled). If you are going to enable TFTP make sure that the –s option is on the line. This specifies that the /tftpboot directory, in this case, is the home directory for TFTP access.

b) Use TFTP only when it is absolutely necessary.

c) Dedicate a system to be a TFTP/BOOTP server.

d) For this example you can also use shadow passwords for the system. When

shadow passwords are enabled, the password information is no longer in the /etc/passwd file. It is moved to a file called /etc/shadow. This file is only readable by root. The hacker would need to get the shadow file in order to crack the passwords.

SNMP – UDP port 161

Simple Network Management Protocol (SNMP) was designed to monitor and manage networks. A network administrator can use a network management application to get or set the information for a managed device through its SNMP agent. These bits of information are called Object Identifiers (OID) and are stored in a management information base (MIB). The OIDs can be used to enable/disable ports, gather user information such as IP and MAC addresses, gather historical information from RMON and much more.

The OIDs that were queried below are common on most, if not all devices with an SNMP agent. With these four OIDs it is possible to determine the software version, possible contact names, what type of interface is used and if those interfaces are up/down. Note that ifAdminStatus not only lets you know if an interface is up or down, but you can also set this OID (ie. You can disable the interface).

Agent Name: test IP Address: 172.26.1.1

Object	Instance	Type	Value
sysDescr	0	OCTET STRING	Hardware: x86 Family 6 Model 5 Stepping 2 AT/AT COMPATIBLE - Software: Windows NT Version 4.0 (BuildNumber 1381 Uniprocessor Free)
sysContact	0	OCTET STRING	John Smith
ifDescr	1	OCTET STRING	MS TCP Loopback interface
ifDescr	2	OCTET STRING	3Com 3C90x Ethernet Adapter
ifAdminStatus	1	INTEGER	up (1)
ifAdminStatus	2	INTEGER	up (1)

SNMPv1 does have a simple authentication mechanism. The 'security' comes in the form of community names. Community names act as form of passwords when contacting an SNMP agent. There is usually a read-only and a read-write community name. Read-only community names will only allow the network management application to retrieve data. Most vendors will use 'public' as the read-only community name. Read-write allows for both retrieval and setting. The default community name that is setup for read-write by most vendors is 'private'. The major problem is that these community names are transmitted over the network in clear text. This can easily be seen in a capture of a SNMPv1 packet:

```
ADDR  HEX                                                ASCII
0000: 00 c0 4f 24 5f 0d 00 e0 63 34 42 8e 08 00 45 00 | ..O$_...c4B...E.
0010: 00 8d 1a d0 00 00 3b 11 61 75 ac 1a 41 0a ac 1a | ......;.au..A...
0020: 8f 69 00 a1 04 02 00 79 78 66 30 82 00 6d 02 01 | .i.....yxf0..m..
0030: 00 04 06 70 75 62 6c 69 63 a2 82 00 5e 02 04 1a | ...public...^...
0040: 7a f5 e8 02 01 00 02 01 00 30 82 00 4e 30 82 00 | z.è......0..N0..
0050: 4a 06 08 2b 06 01 02 01 01 01 00 04 3e 4c 69 6e | J..+.......>Lin
0060: 75 78 20 74 75 6e 6e 65 6c 73 65 72 76 65 72 20 | ux tunnelserver
0070: 32 2e 32 2e 31 36 20 23 38 20 54 68 75 20 53 65 | 2.2.16 #8 Thu Se
0080: 70 20 31 34 20 31 32 3a 35 33 3a 30 37 20 45 44 | p 14 12:53:07 ED
0090: 54 20 32 30 30 30 20 69 35 38 36                | T 2000 i586
```

SNMP section decoded:

```
SNMP: ----- Simple Network Management Protocol (Version 1) -----
      SNMP:
      SNMP: Version      = 1
      SNMP: Community    = public
      SNMP: Command      = Get response
      SNMP: Request ID   = 444265960
      SNMP: Error status = 0 (No error)
      SNMP: Error index  = 0
      SNMP:
      SNMP: Object = {1.3.6.1.2.1.1.1.0} (sysDescr.0)
      SNMP: Value  = Linux tunnelserver 2.2.16 #8 Thu Sep 14 12:53:07 EDT 2000 i586
      SNMP:
```

As one can see, the community name is 'public'.

The networking community came to realize this was a security risk. Work was begun on SNMPv2 (there were actually three versions of SNMPv2) to address the security issues and provide other enhancements. The current version is SNMPv3. It expanded upon the best aspects of the three versions of SNMPv2.

The following suggestions may help reduce the risk of SNMP being used against you.

a) Setup network management stations to poll devices using ping instead of SNMP. This should help limit the amount of SNMP traffic on the network and provide less chances for a community name to be obtained. If you have to use SNMP, use the read-only community name.

b) Try to initially setup a network device in a controlled environment before deploying. This way you can use the read-write community names, if necessary, without fear

that someone will be able to obtain it. Once the device is deployed, try to use the read-only community name (granted there will be times when the read-write community name must be used)

c) If you are using SNMPv1, you should change your community names as frequently as possible. This is actually good practice for any password.

d) Check with the vendors of network devices that support SNMP to see if/when they will support SNMPv3. Update if possible. As a note, SNMPV2 is not used by many vendors. Most vendor are skipping directly to version 3.

Conclusion:

This paper only touches the tip of the iceberg when it comes to information gathering. It should show anyone how easy it is for a hacker to use a few simple tools to gain the data that is needed to potentially break into a system. We have to limit the amount of information that a hacker will obtain. Use the same tools that a hacker would. If you see the same holes in your system you can work to patch them. Upgrade to the latest version of software that is used for a device. Don't give them the chance to break in by using well documented problems in old software. If you need assistance, by all means, hire a professional. In the long run it will probably save you money.

References:

Naugle, Matthew. "Network Protocols Signature Edition." Printed in USA: McGraw-Hill Companies, The, April 1998: 522-562 -- ISBN 0-07-046603-3.

Chappell, Laura. "You're Being Watched – Cyber-crime Scans." Novell Connection, March 2001 (2001): 20-31.

Shipley, Greg. "Tools From the Underground." Network Computing. 29 May 2000. URL: http://www.networkcomputing.com/1110/1110ws1.html (29 March 2001).

Harari, Eddie. "Post-Installation Security Procedures." Linux Journal issue #68. December 1999. URL: http://noframes.linuxjournal.com/lj-issues/issue68/3554.html (3 April 2001).

Lindberg, G., "Anti-Spam Recommendations for SMTP MTAs." Network Working Group Request for Comments: 2505. February 1999.
URL: http://www.ietf.org/rfc/rfc2505.txt?number=2505 (4 April 2001).

Perkins, David T. "SNMP Versions." The Simple Times -- The Quarterly Newsletter of SNMP Technology, Comment, and Events. Volume 5, Number 1, December 1997.
URL: http://www.simple-times.org/pub/simple-times/issues/5-1.html#alternative

(4 April 2001).

Gibson, Steve. "Internet Connection Security for Windows Users." 2001 URL: http://grc.com/lt/howtouse.htm (6 April 2001).

http://www.isi.edu/in-notes/iana/assignments/port-numbers

http://www.insecure.org/nmap/index.html