



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Vulnerabilities In TCP And UDP Ports

Robert Howard Davis

This paper is related to GSEC Version 1.4b – option 1

July 10, 2003

Abstract

This paper concerns network vulnerabilities and defenses, with a focus on TCP and UDP ports.

The make-up of ports will be presented, and the banner message that is sometimes advertised on these ports. Weaknesses in network security will be explored, and also how port scanners can be used to exploit some of these weaknesses.

After a discussion of the risks that are associated with ports, a defense of the network will be presented. A successful defense of the network will involve a series of defensive measures. Everything from preventing access to the port, to protecting the network after a port has been breached. Banners and port scanners will be shown to be useful tools in attacking as well as defending a network device.

Ports Defined

In TCP/IP, a *port* is the mechanism that allows a computer to simultaneously support multiple communication sessions with computers and programs on the network. A port is basically a refinement of an IP address; a computer that receives a packet from the network can further refine the destination of the packet by using a unique port number that is determined when the connection is established.¹

A port is essentially a way for 2 devices to connect using a specific protocol. Each device has an IP address, but this only identifies the device on the network. The port is used to tell each device what kind of a connection will be made. The source and destination port numbers are in the first 31 bits of the packet header. There are three categories of ports. The Internet Assigned Numbers Authority (IANA) developed these categories.

- Numbers 1 through 1023 are Well Known Ports
- Numbers 1024 through 49151 are Registered Ports
- Numbers 49152 through 65535 are Dynamic Ports

Well known ports are described by IANA as ports that generally “can only be used by system (or root) processes or by programs executed by privileged users.”² The ports in this range 0-1023 are registered with IANA. As well as being registered, these ports are also assigned a specific network protocol. Well known ports are usually used to make some kind of network connection using a specific protocol. For instance, the standard telnet port is 23. One device issues a command to make a telnet connection to another device. The command will

identify the protocol to be used. If a port number is not specified in the command, the IP stack of the operating system, will usually assume the Well Known Port number of 23. This number will be put into the packet that is sent to the other device. The device receiving this packet will see that the destination port is 23. At this point the operating system of the device will check the port number and should identify port 23 as the Well Known Port for telnet. In fact this may not happen, as I will explain in the next example.

Actually any port number can be used to make the telnet connection. The only requirement is that both devices are expecting the same protocol on the same port.

My company uses Xyplex terminal servers. The terminal server is a network device that allows asynchronous terminals to make network connections. It has nothing to do with Microsoft terminal services. The Xyplex terminal server expects telnet connections to come on port 2000. A telnet connection to a Xyplex terminal server using port 23 will not work.

I can make a telnet connection from a PC's DOS prompt, to a Xyplex terminal server by issuing the command "telnet" followed by the IP address of the terminal server, a space and the number 2000.

The telnet connection is made from the PC, because the command specified the port 2000. The connection is made at the terminal server because the Terminal server operating system identifies port 2000 as a telnet port.

If I type the telnet command without specifying a port number, my PC will look in the IP stack of the operating system, see that telnet uses port 23, and put that port number into the connection request. The terminal server will not recognize this as a telnet request and the connection will fail.

Examples of Common, Well Know Port Numbers Would Include:

- 21 FTP
- 23 Telnet
- 25 SMTP
- 80 HTTP

Registered Ports

IANA defines registered ports as ports that "can be used by ordinary user processes or programs executed by ordinary users."² These ports are available to any program wanting to use a specific port. If you send a packet to a network device, using a registered port, the operating system of that device should not decide that a registered port is dedicated to any specific protocol. IANA registers the port numbers in this range, but no common network protocol is assigned to them.

Dynamic Ports

Dynamic ports are “unassigned and unregistered ports for private applications, client-side processes, or other processes that dynamically allocate port numbers”.¹ IANA has this advice regarding the use of unassigned ports: “UNASSIGNED PORT NUMBERS SHOULD NOT BE USED”.² It is recommended that an application be sent to IANA to have a port number assigned. The port number should not be used until it is assigned.

Port Banners

Port banners are text descriptions that may appear when a port is accessed. The port banner message may contain information about the protocol used on the port, the operating system or application. Then again, the port may not have a banner at all.

Port Scanners

Ports scanners will try to determine a number of things on the network.

- If an IP address is in use
- If a port connected to an IP address is open.
- The application that is accessing a port. The application could be a PC, file server, firewall, network monitor.
- The version number of the application

In addition, some scanners will show port banners, and also try to attach software using known weaknesses.

Port scanning techniques include:

- Identifying IP addresses is accomplished with “ICMP echo scanning”³ In other words a ping. This is not technically a port scan, as it does not make a port connection.
- “TCP connect()”³ This will attempt to make an actual connection to every port on the device. The machine that was scanned will log these connections.
- “Strobe”⁴ This is a port scan of specific ports.
- “Stealth Scan”⁴ This is a port scan that is not logged. Port scanners use a variety of methods to accomplish this. The scan might use “TCP FIN scanning”⁴ This uses a FIN packet, which will go through most firewalls. The FIN packet is used to terminate a TCP connection. “TCP implements a graceful end by sending a FIN packet followed by an ACK FIN packet from the receiver.”⁵

Some port scanners can be downloaded at no charge from the Internet while others with greater capabilities can be purchased. Port scanners that are freeware include:

- Nmap
- Fscan

- SuperScan
- Nessus
- Wotsweb

The IIS Internet scanner is an example a scanner that can be purchased. When you run the port scanner on an IP address you will see one of three things:

- If a port is not open, you will see no information about that port.
- You may see a message such as “23 telnet”. This is not a port banner. The port scanner sees that port 23 is open and since it is in the Well Known port range, the scanner assumes that port 23 is used for telnet. So this message comes from the port scanner itself.
- The third possibility is a port banner message. When I run the SuperScan port scanner on one of my Unix servers, I see that port 25 (Simple Mail Transfer) is open. When I click on port 25 I am able to see the Fully Qualified Domain Name, the operating system and version number.

Why Port Banners Involve Risk.

The risk in port banners involves the advertisement of operating system, applications, and version numbers. As an example, when I run a port scanner on one of the file servers on my network, the banner on the FTP port will show the operating system and version number. A hacker can use a port scanner to scan IP addresses and network port numbers. If the ports scanned have a banner, they are put into a database. This process can be automated. When a software company identifies a weakness in one of its applications, an announcement is made. The hacker will see the announcement. It is then just a matter of checking the database for the operating system, application and version number. The hacker can then try to attack the network.

Protecting the Network

Up until now the focus has been on TCP and UDP ports. Network defense, however, requires an examination of all access to equipment and data. “Defense In Depth”⁶ describes this process. This Defense In Depth involves security measures at several points in the network. It also involves keeping system and application versions up to date. Attacks can be made at any weakness in the network. The Nimda worm is one software that attacks several areas of the network “through vulnerabilities in IIS web servers, through infected attachments to e-mail, through default Windows disk shares, and through previously infected machines”.⁶ There is no way to insure that a network will never be compromised. The idea is to make an attack more difficult.

Lock the Door

Physical access to the data is very important. Most successful network attacks

come from the inside. Many break-ins involve putting a boot disk or CD into a server then powering it up with admin privilege. This could happen in an office or home, with a desktop PC, as easily as in a computer room with a file server. The truth is that information can be lost or compromised without breaking into a PC. When people leave an office, most PC's are left on. If there is no automatic locking device such as a screen saver password, anyone with access to the office can walk in and access the PC.

Strengthen Passwords

Password security is a special problem. On the one hand, you want passwords to be long and obscure so that they will be difficult to guess. On the other hand, if you can't remember the password you will not have access to the computer. Forgotten passwords can be very time consuming on a company-wide basis. Technically, forgetting a password is not a denial-of-service attack. The CERT Coordination Center defines a denial-of-service attack as "characterized by an explicit attempt by attackers to prevent legitimate users of a service from using that service".⁷ Forgetting a password may not be intentional, but the help desk in my company spends more time on lost passwords than on any other request. This is much more time than has been lost to any virus that was brought in on a floppy disk or came in on an E-mail. For that matter, it is much more time than has been spent in the prevention of viruses at my company.

When most people are given a password that they can't remember, they will write it on a piece of paper and tape it to the PC. Now of course, anyone with access to the PC has access to the password and network resources.

If given the ability to choose a password, most people will choose the same passwords. There are lists of common passwords that are used with password cracking software, to guess passwords. An example of commonly chosen words could be "love, sex, god, secret, default, unknown, aaa, abc, academia, academic, access, ada, admin, aerobics, airplane, albany, alf".⁸ Some people believe that if a password is not a real word that it will be more difficult to guess.

Here is a list of character strings that are commonly chosen as passwords
"!@#\$\$% !@#\$\$%^ !@#\$\$%^& !@#\$\$%^&* 000000 00000000 0007 007
007007 0246 0249 1022 10sne1 111111 121212 1225 123"⁹

Another solution is to have a "common" password. Most people who work with computers, wind up with dozens of passwords. If you use the same password for all of your accounts, you will only have to remember one password. The drawback of course is that if anyone learns the one password, they will, have access to all of the accounts.

I have found several methods for choosing strong passwords that are easily remembered.

- "Choose a line or two from a song or poem, and use the first letter of each word."¹⁰ An example would be "Seventy-six trombones led the big parade", This could become "sstltbp" The password would be stronger by adding punctuation and making some letters upper and lower case. There is a fine line here, as the more you manipulate the password, the

harder it becomes to remember. The revised password could be "SsTITbP!"

- Another procedure would be to create pronounceable non-words "Alternate between one consonant and one or two vowels, up to eight characters."¹⁰ An example could be "bafedone".¹⁰
- A third method would have you "Choose two short words and concatenate them together with a punctuation character between them." For example: ``dog;rain," ``book+mug," ``kid?goat""¹⁰

With any of these schemes you should require a minimum number of characters in the password, depending on the operating system. The makeup of the password should include upper and lower case letters and numbers. The password should also include non-alphanumeric characters.

Run Anti-virus Software at the Desktop

Many viruses are transported from PC to PC on a floppy disk or through E-mail. Anti-virus software on the PC can catch the virus before it causes damage. At work we use Norton. At home I use AVG. They both work as long as the software is updated on a regular basis.

Don't Open SPAM

A virus requires user cooperation in order to replicate. If you don't open E-mail, the virus can't launch an attack against you. I use Outlook, which has a preview feature. This allows the text contents of an E-mail to be viewed. If it is unclear that the E-mail is SPAM, the message can be reviewed with this feature.

Run a Firewall on the PC

At home I use Zone Alarm. This is an exit filtering firewall. No application on my PC can access the Internet without my approval. If a hacker is able to get software on the PC, Zone Alarm should stop it from uploading information found on the PC, downloading something to damage the software, or spreading the software to some other PC via the Internet.

When I installed Zone Alarm, I set it to prompt me for every access attempt. This is a little time consuming, but not as bad as it sounds. One problem with granting access on a program-by program basis is that it can be difficult to tell what application a program is connected to. As I became familiar with which application was launching a program, I set Zone Alarm to always allow specific programs Internet access. For instance, if you click on the Internet Explorer icon, you want the program to be able to connect to the Internet. Once you identify a program that should always be allowed Internet access, Zone Alarm will not prompt you to grant access for that program again.

Install a Firewall On the Network

A firewall on the network will allow you to filter packets by protocol port. You can also exclude specific IP addresses, or only allow specific addresses. This should be the first line of network defense.

Create a DMZ

Use one of the Ethernet ports on the firewall to isolate servers that access the Internet. Since you expect attacks from the Internet, isolate these servers from the rest of the network. These servers might include Citrix servers, E-mail, web servers and routers from other companies. All of these devices are accessible to the Internet, or an outside entity. If one of these servers is compromised, it will be on the outside of the network. The application servers and the network desktops should still be in a safe place.

Run anti-virus software on the network.

Since one of the methods used to transport viruses is E-mail, making an anti-virus server the MX record should stop inbound viruses before they can get to the E-mail system.

Use the Routers in the Network to Filter by Protocol Port

The firewall is protecting the network from the outside. The routers in the WAN can protect the individual networks. Most routers will need access to an application server in another network. Internet or web server access through port 80 may also be required. All other ports and networks should be denied.

Install Software Upgrades and Service Packs

If firewall software is not up to date, the network may be wide open from the outside. This might also open up any DMZ on the network. Most hackers rely on software vendors to announce weaknesses in software. The announcements usually include information about a patch or fix for the newly discovered weakness. If software patches and version upgrades are up to date, you will be one step ahead of most hackers.

Remove Unneeded Applications From the network

This may seem like strange advice. Why would you have unneeded applications on the network? Unless you have looked for them, you are likely to find formerly used applications and applications which were automatically installed during the operating system installation. If IIS web server is unintentionally installed, port 80 may be open to anyone who wants to make a connection.

Unknown applications running on a server have a greater probability of compromising the network. This is because these applications are not likely to be upgraded or patched. Even if the unused application appears to be safe, it is a good rule of thumb to remove unneeded applications. You never know what security hole may be found in the future. The unwanted applications are also a distraction. It is easier to keep track of what applications should be running on the network equipment if you don't have to look at a lot of clutter.

Another benefit of looking for unused applications is that you may also find something that was installed by a hacker. One way to find unused applications is to look for port banners, which advertise the application. If a hacker can use this information to his advantage, why not use it to strengthen the network? Get out the port scanner. At this point you might want to use one that only uses low level scans, and was not designed to exploit software weaknesses. I use SuperScan and have never seen it unload the software that it was scanning.

Know What Ports are Open

Besides finding previously unknown applications, it is a good idea to know what ports are open on each network device in the network. It is a good idea to run a port scanner on all of the network devices on a periodic basis. This could be once a week. You can compare what is open and advertised with the last port scan. If the network is compromised, you will be likely to see a change. In fact, you will see things change even if the network is not compromised. My network at work is constantly changing.

Suspicious Port Banner Message - Port 1025 Network Blackjack

When I first started using a port scanner I noticed something odd on one of our SQL servers. Port 1025 was advertising the Banner "Network Blackjack". I thought that some malicious software had been attached to the server. When I looked at a list of common port names I expected this port to have a normal looking description. But the common name for port 1025 is in fact "Network Blackjack". After a search of the Internet, I discovered that this is a legitimate port.¹¹ It also revealed that I was not the only one who was alarmed when I first saw the banner. There were numerous discussion groups talking about this banner. No one seems to know where the name came from, but 1025 is the first port number after 1024. This makes it often selected by many applications. Six months ago, when I first looked on the Internet for information about this banner, I found many discussion groups talking about it. When I made the same search recently, I found that most of the hits on an Internet search of "Network Blackjack PORT 1025" will take you to actual gambling sites¹². If denial of service includes the time spent running down the information on port banners, then the gambling sites could be viewed as mounting a kind of attack using the search phrase "network blackjack port 1025" in their web page.

Read and Save Log Files

Many network devices will log information about changes to the operating system and access to the device. File servers and routers will log information to a file. You may have to turn the logging on. You should also change the default location to a place with enough disk space to save everything that is logged. At periodic intervals, maybe the same day you run the port scans, you can archive the log files and compare them with the file from the last week.

Removing TCP and UDP Port Banners

The ability to remove port banners is dependant on the operating system, the version of the operating system and the port in question. Of course if the port banner cannot be removed, one work-around might be to change the port number. This could be problematic on a web server. If you are doing business on the Internet, you want security, but you also want people to be able to use the service.

Most of the file and web servers where I work use Microsoft software.

You can see what banners are being advertised on a Microsoft server by issuing a command at the DOS prompt, or clicking on start, then run¹³. Type telnet, a space, and the "computer name" of the device, another space and a port number. You may have to press the return button again. If there is a banner associated with this port you should see it now. You will have to do this for each port number. An alternate method would be to use a port scanner. Put in the IP address of the computer, and select all ports.

To remove the port banners on an IIS web server you will need to download the IISlockdown tool. This is free from Microsoft and can be obtained from www.microsoft.com/downloads. The actual software needed to remove banners is URLScan. This software is included in the IISlockdown tool. Microsoft recommends the installation of this tool, whether or not you decide to delete port banners. "URLScan is also very helpful in protecting the web server from present and future vulnerabilities like Code Red / Nimda and is highly recommended"¹³

Remove Telnet Banners from Port 80 (HTTP) on the IIS Server.

The IIS server has a port banner on port 80 (HTTP) that may identify the operating system and the version of the operating system. This banner can be removed if the IIS server is version 4 or above. To remove the telnet banner, use the IISlockdown tool to install URLscan. Then stop the IISAdmin service, and edit the file "%systemroot%\System32\Inetsrv\Urlscan\Urlscan.ini".¹⁴ The contents of the file will be "RemoveServerHeader=0"¹⁴. Edit the file and change the "0" to a "1". The IIS server can then be restarted with the IISRESET command.

Change POP and IMAX Port Banners on the Windows Exchange 2000 Server

The default information in the banner of an Exchange server 2000 will identify the application as a Microsoft Exchange server and give the version number. There are 2 banners for each of these 2 protocols. The first is a connection banner. The second is a disconnection banner.

One command is used to change each of these 4 banners. The command must identify the banner to be changed with a metabase key number that identifies one of the 4 banners. The POP or IMAX port must also be specified in the command. Therefore, the command must be executed 4 times to change all 4 banners.

The procedure is to stop the service and issue the command that will change the value in the server metabase. The service can then be restarted.

The command used to change the banner on the IMAP4 connection string would be "smtpmd SET -path imap4svc/1 -dtype STRING -prop 49884 -value "<the_new_connection_string>"¹⁵

The number 49884 is the metabase key for the IMAP4 connection string.

The 4 possibilities are:

- POP3 Connection String 41661
- POP3 Disconnection String 41662
- IMAP4 Connection String 49884
- IMAP4 Disconnection String 49885

Port banners on Versions of the Exchange server prior to Exchange 2000 cannot be changed.

Changing the SMTP Port Banner From the Microsoft Exchange 2000 Server

The port banner for SMTP on a 2000 server can be changed, however Microsoft issues the following warning:

WARNING: If you use the ADSI Edit snap-in, the LDP utility, or any other LDAP version 3 client, and you incorrectly modify the attributes of Active Directory objects, you can cause serious problems. These problems may require you to reinstall Microsoft Windows 2000 Server, Microsoft Exchange 2000 Server, or both. Microsoft cannot guarantee that problems that occur if you incorrectly modify Active Directory object attributes can be solved. Modify these attributes at your own risk.¹⁶

The Windows 2000 SMTP banner will show the fully qualified domain name, identify the application as Microsoft ESMTP MAIL, the version number, and the date. While the banner cannot be removed, it can be altered so that it only shows the fully qualified domain name and the date.

Change the banner by modifying the file Lm\Smtpsvc\virtual server number".

Use a metabase editing tool to do this. Microsoft suggests MetaEdit. "The Metabase Editor (MetaEdit) is a tool that provides similar functionality to the Windows NT Registry Editor. Using MetaEdit, you can browse and modify attributes in the Metabase. Note that in using MetaEdit, you can make changes that may damage your IIS configuration."¹⁷

Select Edit, New, and then String. The entry in the ID box should be "Other". Now put 36907 to the right of the ID box. Type the new banner string into the Value box. Now restart the SMTP service.

Microsoft FTP and NNTP Port Banners Cannot Be Removed.

There is no way to remove FTP and NNTP port banners from any Windows operating system. A workaround could be to use a port in the Registered Port range of 1024 through 49151.

Conclusion

TCP and UDP ports can be used to attack network devices. As vulnerabilities related to these ports are discovered, software companies will develop and release code to correct the weaknesses. The information in these releases can be used to attack network devices that have not yet installed the corrective software.

Network devices can be protected by the timely installation of software updates. Successful attacks on TCP and UDP ports may also be prevented by a Defense in Depth. This involves a number of defenses that will deter an attacker. If physical access to a device is not available, an attack may be prevented. If a password cannot be guessed, an attack may fail, even though a connection has been made to a network device.

No one defense can guarantee that a network device will be safe from attack. The best defense of network devices is a series of measures, which will make attacks as difficult as possible.

© SANS Institute 2000 - 2005. All rights reserved. SANS Institute retains full rights.

[1]

Microsoft TechNet "Appendix B - Port Reference for MS TCP/IP."

Well known services are defined by RFC 1060 July 1992

URL:http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/winntas/reskit/net/port_nts.asp

[2]

IANA Port Numbers May 20, 2003

URL:<http://www.iana.org/assignments/port-numbers>

[3]

The Art of Port Scanning by Fyodor

Last significant update: Sat Sep 6 03:24:53 GMT 1997

URL:http://www.insecure.org/nmap/nmap_doc.html

[4]

Webopedia Port scanning

Types of port scans - February 2, 2002

URL:http://www.webopedia.com/TERM/P/port_scanning.html

[5]

Internet Connection Software Page...

TRANSPORT CONTROL PROTOCOL (TCP)

URL:<http://www.netfor2.com/tcp.htm>

[6]

Defense in Depth Benefits

URL:<http://securityresponse.symantec.com/avcenter/security/Content/security.articles/defense.in.depth.html>

[7]

CERT Coordination Center

Denial of Service Attacks June 4, 2001

URL:http://www.cert.org/tech_tips/denial_of_service.html#1

[8]

PROTECT YOUR PRIVACY ON THE INTERNET: PASSWORDS 2002

URL:[HTTP://www.taciroglu.com/p/p.htm](http://www.taciroglu.com/p/p.htm)

[9]

GeodSoft Website Consulting 2000

URL:<http://geodsoft.com/howto/password/common.htm>

[10]

Excerpts from IMPROVING THE SECURITY OF YOUR UNIX SYSTEM

David A. Curry, Systems Programmer
Information and Telecommunications Sciences and Technology Division
URL:<http://www.alw.nih.gov/Security/Docs/passwd.html>

[11]
TheoryGroup Re: UDP port 1025/6 July 23,2000
URL:<http://www.theorygroup.com/Archive/Unisog/2000/msg00904.html>

[12]
City Club Casino July 2003
URL:http://www.101-best-online-casinos.com/s_port_1025_blackjack_1.shtml

[13]
Microsoft FAQ
(11) IIS INTERNET INFORMATION SERVICES [WEB SERVER, FTP SERVER, ETC]
SEEING THE BANNERS ON YOUR COMPUTER:
URL:<http://securityadmin.info/faq.asp#iis>

[14]
Microsoft Knowledge Base Article - 317741
Mask IIS Version Information from Network Trace and Telnet May 20 2003
URL:<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q317741>

[15]
Microsoft Knowledge Base Article - 303513
XCON: How to Modify the POP or IMAP Banner June 5,2003
URL:<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q303513>

[16]
Microsoft Knowledge Base Article - 281224
XCON: How to Modify the SMTP Banner June 4, 2003
<http://support.microsoft.com/default.aspx?scid=kb;EN-US;281224>
4

[17]
MetaEdit - edit the IIS Metabase October 5, 2001
<http://www.webattack.com/get/metaedit.sht>