



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.

GIAC Firewall and perimeter Protection Practical by David Leach

Network Configuration

Our test network consists of one subnet 95.5.6.0/24 which consists of 254 valid hosts, the IP address range is 95.5.6.0 -- 95.5.6.255. Our network consists of a Cisco 3600 Series Router running IOS 11.3. Our firewall is a Mandrake 7.1 Linux machine running ipchains. A 3Com switch 1000 12 port switch is used in our screened network. We are using a 3Com Corebuilder 5000 Switch as the backbone switch on our trusted network. All PC's are running Windows NT 4.0 Workstation SP6a. Four file and 2 print servers are running Windows NT 4.0 Server SP6a. The database server is running Windows NT 4.0 Server SP6a and MS-SQL server 7.0 SP2. Three DNS servers are running Redhat Linux 6.1 and Bind version 8.2.2-P5 . The web server is running Redhat Linux 6.1 and Apache 1.3.11. The internal SMTP server is running Windows NT 4.0 Server SP6a and MS-Exchange Server 5.5 SP3. The external SMTP server is running Redhat Linux 6.1 and sendmail 8.11.0.

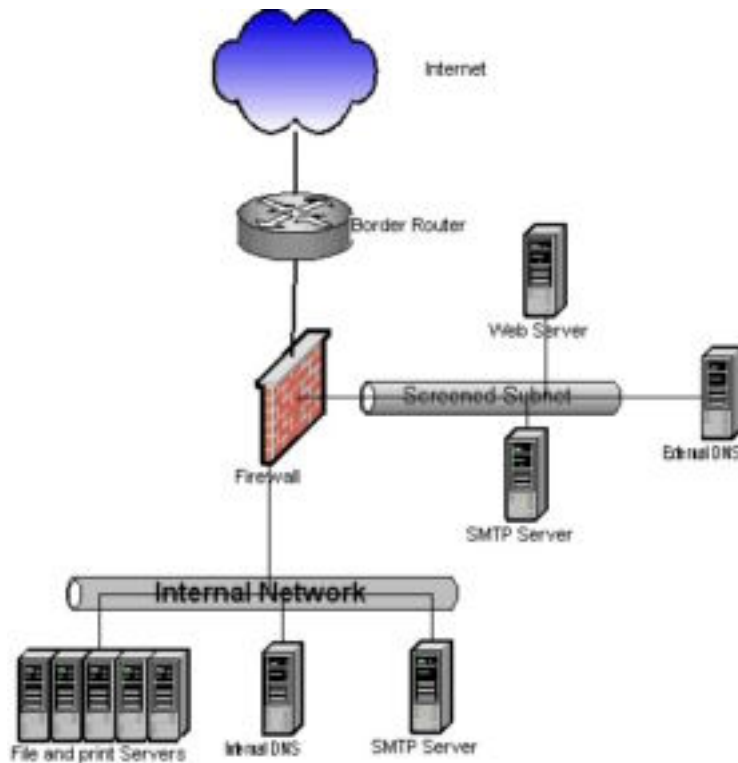
The external (s0) interface of the Cisco Router is using IP unnumbered, our internal (e0) interface IP address is 95.5.6.1. We are using the IP featureset and have several extended access lists installed on this machine. Our Internet connection is a T1 providing 1.54Mbps bandwidth to our ISP.

The firewall is configured with 3 10/100 network interface cards (NICs). The external NIC (eth2) is connected to the Cisco Router via a crossover cable and its IP address is 95.5.6.2. The DMZ NIC (eth1) is connected to the 3Com Switch 1000 and its IP address is 95.5.6.3. The internal NIC (eth0) is connected to the 3Com Corebuilder 5000.

Our screened network contains 3 servers, the primary external DNS server (IP address 95.5.6.10, the web server (IP address 95.5.6.20) and the SMTP server (IP address 95.6.5.25). The secondary external DNS is hosted by our ISP (IP address 73.5.15.10).

The exchange server (IP address 95.5.6.225) on our trusted network handles internal e-mail, calanders, and contact lists and sends e-mail destine for outside networks to the external SMTP server. The internal DNS servers (Primary DNS IP address 95.5.6.210, Secondary DNS IP address 95.5.6.211) handle DNS requests from internal hosts only.

© SANS Institute 2000 - 2002



Network Security Policy

Our security policy will allow http traffic to the Internet from our users, SMTP communication between our internal exchange server and our external sendmail server and DNS queries from our DNS servers. We also block all spoofed addresses from entering and exiting our network, source routed packets, logon services, RPC, NFS, NetBios, X-Windows, Naming Services, small services, TFTP, finger, NNTP, NTP, LPD, syslog, SNMP, BGP, SOCKS and certain ICMP services.

We will create a filter ruleset on our border router and implement a linux firewall with ipchains. The following are the commands needed to implement these rules on the appropriate devices.

(items in italics are provided for informational purposes and should not be entered at the command line on you router)

Login to the router and enter enable/privileged mode and enter the following commands one per line:

```
No ip source-route (discards any packet with the source route option enabled)
Access-list 100 deny ip 0.0.0.0 0.255.255.255 any (deny all traffic with a source of 0.x.x.x)
Access-list 100 deny ip 1.0.0.0 0.255.255.255 any (deny all traffic with a source of 1.x.x.x)
Access-list 100 deny ip 10.0.0.0 0.255.255.255 any (deny all traffic with a source of 10.x.x.x)
Access-list 100 deny ip 23.0.0.0 0.255.255.255 any (deny all traffic with a source of 23.x.x.x)
Access-list 100 deny ip 31.0.0.0 0.255.255.255 any (deny all traffic with a source of 31.x.x.x)
Access-list 100 permit ip 65.0.0.0 0.255.255.255 any (Permit all traffic with a source of 65.x.x.x)
Access-list 100 permit ip 66.0.0.0 0.255.255.255 any (Permit all traffic with a source of 66.x.x.x)
Access-list 100 permit ip 95.0.0.0 0.255.255.255 any (Permit all traffic with a source of 95.x.x.x)
Access-list 100 deny ip 65.0.0.0 31.255.255.255 any (deny all traffic with a source between 65.x.x.x and 95.x.x.x)
Access-list 100 deny ip 96.0.0.0 31.255.255.255 any (deny all traffic with a source between 96.x.x.x and 127.x.x.x)
Access-list 100 deny ip 128.66.0.0 0.0.255.255 any (deny all traffic with a source of 128.66.x.x)
Access-list 100 deny ip 128.0.0.0 0.0.255.255 any (deny all traffic with a source of 128.0.x.x)
Access-list 100 deny ip 172.16.0.0 0.15.255.255 any (deny all traffic with a source between 172.16.x.x and 172.32.x.x)
Access-list 100 deny ip 191.255.0.0 0.0.255.255 any (deny all traffic with a source of 191.255.x.x)
Access-list 100 deny ip 192.0.0.0 0.0.0.255 any (deny all traffic with a source of 191.0.0.x)
Access-list 100 deny ip 192.0.1.0 0.0.0.255 any (deny all traffic with a source of 191.0.1.x)
Access-list 100 deny ip 192.0.2.0 0.0.0.255 any (deny all traffic with a source of 191.0.2.x)
Access-list 100 deny ip 192.168.0.0 0.0.255.255 any (deny all traffic with a source of 192.168.x.x)
```

Access-list 100 deny ip 197.0.0.0 0.255.255.255 any (deny all traffic with a source of 197.x.x.x)
Access-list 100 deny ip 201.0.0.0 0.255.255.255 any (deny all traffic with a source of 201.x.x.x)
Access-list 100 deny ip 223.255.255.0 0.0.0.255 any (deny all traffic with a source of 223.255.255.x)
Access-list 100 deny ip 224.0.0.0 31.255.255.255 any (deny all traffic with a source between 224.x.x.x and 239.x.x.x)
Access-list 100 deny ip 240.0.0.0 15.255.255.255 any (deny all traffic with a source between 240.0.x.x and 255.x.x.x)
Access-list 100 deny ip 95.5.6.0 0.0.0.255 any (deny all traffic with a source of 95.5.6.x)
Access-list 100 permit ip any any (Permit all other traffic)
Config t (enter configure terminal mode)
Int s0 (the external interface to apply the list to)
Ip access-group 100 in (apply/group the access list 100 inbound to this interface)
Exit (exit the interface configuration mode)
Exit (exit the configure terminal mode)
Write (write the running/changed configuration to NVRAM)

Create a Linux machine with three interfaces on it according to the above description. Create a file called chains (the name doesn't really matter) with on that machine which is identical to the following text between the lines:

```
:login -
:web -
:rpc -
:netbios -
:misc -
:XWin -
:name -
:mail -
:icmp -
:small -
:input ACCEPT -j login
:input ACCEPT -j web
:input ACCEPT -j rpc
:input ACCEPT -j netbios
:input ACCEPT -j misc
:input ACCEPT -j XWin
:input ACCEPT -j name
:input ACCEPT -j mail
:input ACCEPT -j icmp
:input ACCEPT -j small
:output ACCEPT
:forward ACCEPT
# Login Services
-A login -s 0/0 -d 95.5.6.0/24 20:23 -i eth2 -p 6 -j DENY -l
-A login -s 0/0 -d 95.5.6.0/24 512:514 -i eth2 -p 6 -j DENY -l
#
# Addressed in the NetBios section
#-A login -s 0/0 -d 95.5.6.0/24 139 -i eth2 -p 6 -j DENY -l
#
#RPC and NFS Services
-A rpc -s 0/0 -d 95.5.6.0/24 111 -i eth2 -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 2049 -i eth2 -y -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 4045 -i eth2 -y -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 111 -i eth2 -p 17 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 2049 -i eth2 -p 17 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 4045 -i eth2 -p 17 -j DENY -l
#
#NetBios for NT and Win2K
-A netbios -s 0/0 -d 95.5.6.0/24 135 -i eth2 -p 6 -j DENY -l
-A netbios -s 0/0 -d 95.5.6.0/24 135 -i eth2 -p 17 -j DENY -l
```

```

-A netbios -s 0/0 -d 95.5.6.0/24 137:139 -p 6 -j DENY -l
-A netbios -s 0/0 -d 95.5.6.0/24 137:139 -p 17 -j DENY -l
#
#X Windows Services
-A XWin -s 0/0 -d 95.5.6.0/24 6000:6255 -y -p 6 -j DENY -l
#
#Name Services
-A name -s 0/0 -d 95.5.6.0/24 389 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 389 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.10/32 53 -p 17 -j ACCEPT
#
#Mail Services
-A mail -s 0/0 1024:65535 -d 95.5.6.25/32 25 -p 6 -j ACCEPT
-A mail -s 95.5.6.25/32 1024:65535 -d 0/0 25 -p 6 -j ACCEPT
-A mail -s 95.5.6.25/32 1024:65535 -d 172.16.1.25/32 25 -p 6 -j ACCEPT
-A mail -s 172.16.1.25/32 1024:65535 -d 95.5.6.25/32 25 -p 6 -j ACCEPT
-A mail -s 0/0 -d 95.5.6.0/24 25 -p 6 -j DENY -l
-A mail -s 0/0 -d 95.5.6.0/24 109:110 -p 6 -j DENY -l
-A mail -s 0/0 -d 95.5.6.0/24 143 -p 6 -j DENY -l
#
#Web Services
-A web -s 0/0 -d 95.5.6.20/32 80 -p 6 -j ACCEPT
-A web -s 0/0 -d 95.5.6.20/32 443 -p 6 -j ACCEPT
-A web -s 0/0 -d 95.5.6.0/24 80 -p 6 -j DENY -l
-A web -s 0/0 -d 95.5.6.0/24 8000 -y -p 6 -j DENY -l
-A web -s 0/0 -d 95.5.6.0/24 8080 -y -p 6 -j DENY -l
-A web -s 0/0 -d 95.5.6.0/24 8888 -y -p 6 -j DENY -l
#
#Miscellaneous Services
-A name -s 0/0 -d 95.5.6.0/24 79 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 69 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 119 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 123 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 515 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 161:162 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 161:162 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 179 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 1080 -y -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 514 -p 17 -j DENY -l
#
#Icmp Filter
-A icmp -s 0/0 -d 95.5.6.0/24 --icmp-type echo-request -p 1 -j DENY -l
-A icmp -s 95.5.6.0/24 -d 0/0 --icmp-type echo-reply -p 1 -j DENY -l
-A icmp -s 0/0 -d 0/0 --icmp-type time-exceeded -p 1 -j DENY -l
-A icmp -s 0/0 -d 0/0 --icmp-type destination-unreachable -p 1 -j DENY -l
#
#Small Servers
-A small -s 0/0 -d 0/0 0:20 -p 6 -j DENY -l
-A small -s 0/0 -d 0/0 0:20 -p 17 -j DENY -l
-A small -s 0/0 -d 0/0 37 -p 6 -j DENY -l
-A small -s 0/0 -d 0/0 37 -p 17 -j DENY -l

```

To apply the above rule set you can use the command: `ipchains-restore < chains`

Explanation of blocked ports

Block "spoofed" addresses-- packets coming from outside your company sourced from internal addresses or private (RFC1918 and network 127) addresses. Also block source routed packets.

A spoof

as defined at "whatis.com":

"Spoof" was a game involving trickery and nonsense that was invented by an English comedian, Arthur Roberts, prior to 1884, when it is recorded as having been "revived." Webster's defines the verb to mean (1) to deceive or hoax, and (2) to make good-natured fun of.

On the Internet, "to spoof" can mean:

- 1) To deceive for the purpose of gaining access to someone else's resources (for example, to fake an Internet address so that one looks like a certain kind of Internet user)
- 2) To simulate a communications protocol by a program that is interjected into a normal sequence of processes for the purpose of adding some useful function
- 3) To playfully satirize a Web site.

Possible vulnerabilities associated with spoofed addresses

For our purposes address "spoofing" occurs when an attacker manipulates the source IP address of a packet as to hide their true location. Address spoofing is often used in Denial of Service (DoS) attacks when the attacker is not concerned with the response from the victim host. When "spoofed" addresses are used in an attack the attacker can divert attention from their location to another site, make it more difficult to track and possibly bypass many of the rules in a firewall, IDS or other protection devices. By blocking network addresses that we know can not originate outside of our organization we can reduce an attackers ability to use certain addresses when targeting our systems. These addresses can be the private addresses as outlined in RFC1918,

10.0.0.0 - 10.255.255.255 (10/8 prefix)
172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

Also outlined by the Internet Assigned Number Authority (IANA) are a number of reserved addresses that we will also treat as spoofed addresses.

[No name] (NET-TEST)	IANA	192.0.2.0
IANA (BLACKHOLE3-HST)	BLACKHOLE.ISI.EDU	128.9.64.26
IANA (RESERVED-4)	RESERVED	191.255.0.0
IANA (RESERVED-5)	RESERVED	223.255.255.0
IANA (LOOPBACK)	LOOPBACK	127.0.0.0
IANA (NET-TEST-B)	TEST-B	128.66.0.0
IANA (RESERVED-9)	RESERVED-9	1.0.0.0 - 1.255.255.255
IANA (RESERVED-3)	RESERVED-3	128.0.0.0
IANA (RESERVED-1)	RESERVED-1	0.0.0.0 - 0.255.255.255
IANA (RESERVED-8)	RESERVED-8	96.0.0.0 - 126.255.255.255
IANA (RESERVED-7)	RESERVED-7	67.0.0.0 - 95.255.255.255
IANA (NET-DDN-TC-NET)	RESERVED-23	23.0.0.0 - 23.255.255.255
IANA (RESERVED-12)	RESERVED-31	31.0.0.0 - 31.255.255.255
IANA (NET-ROOT-NS-LAB)	ROOT-NS-LAB	192.0.0.0
IANA (NET-ROOTS-NS-LIVE)	NET-ROOTS-NS-LIVE	192.0.1.0 - 192.0.1.255
IANA (IANA-CBLK-RESERVED)	IANA-CBLK1	192.168.0.0 - 192.168.255.0
IANA (IANA-BBLK-RESERVED)	IANA-BBLK-RESERVED	172.16.0.0 - 172.31.0.0
IANA (NETBLK-RESERVED-13)	RESERVED-13	197.0.0.0 - 197.255.255.255
IANA (NETBLK-RESERVED-14)	RESERVED-14	201.0.0.0 - 201.255.255.255
IANA (RESERVED-6)	RESERVED-10	10.0.0.0 - 10.255.255.255
IANA (RESERVED-2)	RESERVED-192	192.0.0.0 - 192.0.255.255

The addresses assigned to us by our ISP 95.5.6.0/24, or the loopback address 127.0.0.1 and its address space 127.0.0.0/8.

Source routed packets

Defined

Source routed packets are packets which have been generated in such a way that they claim to be from inside the firewall. What this essentially means that the packet was routed to your network in one direction and arrives at the firewall claiming to have arrived from a different route. By doing this the

attacker can bypass some firewall and router rules and gain access to trusted systems. With source routing options on the return traffic will be return along the reverse path of the source route. With this option off the return traffic would not be able to reach the attacker.

Possible vulnerabilities associated with Source routed packets

Source routed packets allow an attacker to send packets which appear from inside the trusted network. This has a similar affect to spoofing in that packets may be allowed to pass through the firewall because they appear to have originated from within the trusted network. The difference is that with Source routed packets it is the routing information that has changed not the IP address so the attacker will receive the responses from the server. This could be used in a number of attacks and fingerprinting tools.

Extended Access Control list syntax on the Cisco Router:

```
access-list access-list-number {permit | deny} protocol source [src-wildcard-mask]
destination [dest-wildcard-mask] [operator operand]
```

Access-list – keyword indicating the following will be an access list.

access-list-number – number assigned to this list. This allows us to create a list instead of single entries

{permit | deny} – whether to allow or deny the traffic to flow on this interface.

Protocol – type of protocol to filter ex. Tcp, icmp, udp

source [src-wildcard-mask] – Source address and mask identifying the source network or host to filter on.

destination [dest-wildcard-mask] – Destination address and mask identifying the destination network or host to filter on

[operator operand] – some operators are equal, grater than, less than not equal. Operand is the port you are acting on.

Example: access-list 100 deny tcp 10.0.0.0 0.255.255.255 any eq 23

This example blocks all traffic from the network 10.0.0.0 directed at tcp port 23 (telnet)

```
Access-list 100 deny ip 0.0.0.0 0.255.255.255 any
Access-list 100 deny ip 1.0.0.0 0.255.255.255 any
Access-list 100 deny ip 10.0.0.0 0.255.255.255 any
Access-list 100 deny ip 23.0.0.0 0.255.255.255 any
Access-list 100 deny ip 31.0.0.0 0.255.255.255 any
Access-list 100 permit ip 65.0.0.0 0.255.255.255 any
Access-list 100 permit ip 66.0.0.0 0.255.255.255 any
Access-list 100 deny ip 65.0.0.0 31.255.255.255 any
Access-list 100 deny ip 96.0.0.0 31.255.255.255 any
Access-list 100 deny ip 127.0.0.0 0.255.255.255 any
Access-list 100 deny ip 128.66.0.0 0.0.255.255 any
Access-list 100 deny ip 128.0.0.0 0.0.255.255 any
Access-list 100 deny ip 172.16.0.0 0.15.255.255 any
Access-list 100 deny ip 191.255.0.0 0.0.255.255 any
Access-list 100 deny ip 192.0.0.0 0.0.0.255 any
Access-list 100 deny ip 192.0.1.0 0.0.0.255 any
Access-list 100 deny ip 192.0.2.0 0.0.0.255 any
Access-list 100 deny ip 192.168.0.0 0.0.255.255 any
Access-list 100 deny ip 197.0.0.0 0.255.255.255 any
Access-list 100 deny ip 201.0.0.0 0.255.255.255 any
```

```
Access-list 100 deny ip 223.255.255.0 0.0.0.255 any
Access-list 100 deny ip 224.0.0.0 31.255.255.255 any
Access-list 100 deny ip 240.0.0.0 15.255.255.255 any
Access-list 100 deny ip 95.5.6.0 0.0.0.255 any
Access-list 100 permit ip any any
```

Login services-- telnet (23/tcp), SSH (22/tcp), FTP (21/tcp), NetBIOS (139/tcp), rlogin et al (512/tcp through 514/tcp)

Telnet

as defined at "whatis.com":

Telnet is the way you can access someone else's computer, assuming they have given you permission. (Such a computer is frequently called a host computer.) More technically, Telnet is a user command and an underlying TCP/IP protocol for accessing remote computers. The Web or HTTP protocol and the FTP protocol allow you to request specific files from remote computers, but not to actually be logged on as a user of that computer. With Telnet, you log on as a regular user with whatever privileges you may have been granted to the specific applications and data on that computer.

A Telnet command request looks like this (the computer name is made-up): telnet the.libraryat.harvard.edu

The result of this request would be an invitation to log on with a userid and a prompt for a password. If accepted, you would be logged on like any user who used this computer every day.

Telnet is most likely to be used by program developers and anyone who has a need to use specific applications or data located at a particular host computer.

Possible vulnerabilities associated with Telnet

All Telnet traffic is sent in clear text including passwords. If an attacker is "sniffing" the target network when a telnet session occurs, the attacker will see everything that is sent between the two networks including usernames, passwords, and files and directories that are listed during the session. This information would provide an attacker everything they need to launch a successful attack.

Consequences associated with blocking Telnet:

Without telnet, administrators may not be able to manage systems across the firewall. Also some login services to UNIX machines will also be blocked.

SSH

as defined at "whatis.com":

Secure Shell (SSH) is a UNIX-based command interface and protocol for securely getting access to a remote computer. It is widely used by network administrators to control Web and other kinds of servers remotely. SSH is actually a suite of three utilities - slogin, ssh, and scp - that are secure versions of the earlier UNIX utilities, rlogin, rsh, and rcp. SSH commands are encrypted and secure in several ways. Both ends of the client/server connection are authenticated using digital certificates and passwords are protected by being encrypted.

SSH uses RSA Laboratory's public key cryptography for both connection and authentication. Encryption algorithms include Blowfish, DES, and IDEA. IDEA is the default.

SSH2, the latest version, is a proposed set of standards from the Internet Engineering Task Force (IETF).

Possible vulnerabilities associated with SSH

One of the things that make SSH so powerful and dangerous at the same time is the ability to tunnel one protocol or application through SSH. What that means is once a SSH communication channel has been established all traffic between the two hosts is encrypted. Most SSH programs have a "port redirect" feature which enables a program, such as Netscape, to send information to the SSH application which will in turn encrypt the traffic and send it to the host which it is communicating. Once on that host it will be decrypted and possibly sent on from there. This would allow an attacker to send malicious code to a server inside a trusted network using the SSH encryption tunnel. This is even more dangerous than just having the attacker send to code to the target server because the communications are encrypted and an IDS sensor will not decode the traffic and identify the problem.

Consequences associated with blocking SSH

Without SSH users and administrators would not be able to communicate through the firewall using SSH encryption and tunneling.

FTP

as defined at "whatis.com":

FTP (File Transfer Protocol), a standard Internet protocol, is the simplest way to exchange files between computers on the Internet. Like the Hypertext Transfer Protocol (HTTP), which transfers displayable Web pages and related files, and the Simple Mail Transfer Protocol (SMTP), which transfers e-mail, FTP is an application protocol that uses the Internet's TCP/IP protocols. FTP is commonly used to transfer Web page files from their creator to the computer that acts as their server for everyone on the Internet. It's also commonly used to download programs and other files to your computer from other servers.

As a user, you can use FTP with a simple command line interface (for example, from the Windows MS-DOS Prompt window) or with a commercial program that offers a graphical user interface. Your Web browser can also make FTP requests to download programs you select from a Web page. Using FTP, you can also update (delete, rename, move, and copy) files at a server. You need to log on to an FTP server. However, publicly available files are easily accessed using anonymous FTP.

Basic FTP support is usually provided as part of a suite of programs that come with TCP/IP. However, any FTP client program with a graphical user interface usually must be downloaded from the company that makes it

Possible vulnerabilities associated with FTP

FTP does not use encryption with its passwords therefore it is possible a username and password can be compromised. There are also some ftp servers which have vulnerabilities that allow system compromise. This could lead to denial of service, alteration of websites, or relay of an attack to another site.

Consequences associated with blocking FTP

Without FTP, customers may not be able to download patches or system updates. If FTP is also blocked outgoing internal users may not be able to download operating system patches or other data files.

NetBIOS

as defined at "whatis.com":

NetBIOS (Network Basic Input/Output System) is a program that allows applications on different computers to communicate within a local area network (LAN). It was created by IBM for its early PC Network, was adopted by Microsoft, and has since become a de facto industry standard. NetBIOS is used in Ethernet, token ring, and Windows NT networks. It does not in itself support a routing mechanism so applications communicating on a wide area network (WAN) must use another "transport mechanism" (such as TCP) rather than or in addition to NetBIOS. NetBIOS frees the application from having to understand the details of the network, including error recovery (in session mode). A NetBIOS request is provided in the form of a Network Control Block (NCB) which, among other things, specifies a message location and the name of a destination.

NetBIOS provides the session and transport services described in the Open Systems Interconnection (OSI) model. However, it does not provide a standard frame or data format for transmission. A standard frame format is provided in the NetBIOS Extended User Interface (NetBEUI).

NetBIOS provides two communication modes: session or datagram. Session mode lets two computers establish a connection for a "conversation," allows larger messages to be handled, and provides error detection and recovery. Datagram mode is "connectionless" (each message is sent independently), messages must be smaller, and the application is responsible for error detection and recovery. Datagram mode also supports the broadcast of a message to every computer on the LAN.

Possible vulnerabilities associated with NetBIOS

NetBIOS can allow for global file sharing over the Internet. Due to its design NetBIOS is not a secure product. There are several vulnerabilities that have been identified with Windows machines that can lead to a denial of service or other erratic behavior of the system.

Consequences associated with blocking NetBIOS

This service is usually associated with local communications and should not be normal traffic on the Internet. Unless corporate policy dictates that this is a necessary service it should be disabled.

Rlogin

as defined at "whatis.com":

Rlogin (remote login) is a UNIX command that allows an authorized user to login to other UNIX machines (hosts) on a network and to interact as if the user were physically at the host computer. Once logged in to the host, the user can do anything that the host has given permission for, such as read, edit, or delete files.

Rlogin is similar to the better known Telnet command. Rlogin is considered useful for simple logins that don't

require a lot of control over the client/host interaction, but is thought to be less useful than Telnet where a lot of customization is desired, for multiple sessions, for connections between very distant terminals or to terminals that are not running UNIX, for that matter, since rlogin can only connect to UNIX hosts. A benefit of rlogin is the ability to use a file called `.rhosts` that resides on the host machine and maintains a list of terminals allowed to login without a secure version of rlogin (`slogin`) was combined with two other UNIX utilities, `ssh` and `scp`, in the Secure Shell suite, an interface and protocol created to replace the earlier utilities.

Rsh

as defined in the man pages:

`rsh` connects to the specified *hostname* and executes the specified *command*. `rsh` copies its standard input to the remote command, the standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, quit, and terminate signals are propagated to the remote command; `rsh` normally terminates when the remote command does. If you omit *command*, instead of executing a single command, `rsh` logs you in on the remote host using `rlogin(1)`. Shell metacharacters which are not quoted are interpreted on the local machine, while quoted metacharacters are interpreted on the remote machine.

Rcopy

as defined in the man pages:

The `rcp` command copies files between machines. Each *filename* or *directory* argument is either a remote file name of the form: *hostname:path* or a local file name (containing no ":" (colon) characters, or "/" (backslash) before any ":" (colon) characters). If a *filename* is not a full path name, it is interpreted relative to your home directory on *hostname*. A *path* on a remote host may be quoted using `\`, `"`, or `'`, so that the metacharacters are interpreted remotely. `rcp` does not prompt for passwords; your current local user name must exist on *hostname* and allow remote command execution by `rsh(1)`. `rcp` handles third party copies, where neither source nor target files are on the current machine. Hostnames may also take the form *username@hostname:filename* to use *username* rather than your current local user name as the user name on the remote host. `rcp` also supports Internet domain addressing of the remote host, so that: *username@host.domain:filename* specifies the username to be used, the hostname, and the domain in which that host resides. File names that are not full path names will be interpreted relative to the home directory of the user named *username*, on the remote host.

Possible vulnerabilities associated with rlogin, rsh, rcopy

The "R services" do not provide any strong authentication. They rely on users and administrators to properly configure the associated files, such as `.rhosts`. When these files are used remote users connect to the local machine with the rights the associated local account has. The only restriction is the IP address. This could lead to a root level compromise of the local machine if the remote machine is compromised or the ip address is spoofed.

Consequences associated with blocking rlogin, rsh, rcopy

Without the ability to use the "R services" administrators must use some other form of connecting to remote hosts. Telnet, `ssh` or even console access only would be appropriate.

Using `ipchains` you can implement these commands to create the restrictions. One can create a file to import these rules into `ipchains`. This file should contain the following lines:

```
-A login -s 0/0 -d 95.5.6.0/24 20:23 -p 6 -j DENY -l
    (block all traffic destined to the 95.5.6.0/27 network on tcp ports 20, 21, 22, 23)
-A login -s 0/0 -d 95.5.6.0/24 512:514 -p 6 -j DENY -l
    (block all traffic destined to the 95.5.6.0/27 network on tcp ports 512, 513, 514)
-A login -s 0/0 -d 95.5.6.0/24 139 -p 6 -j DENY -l
    (block all traffic destined to the 95.5.6.0/27 network on tcp port 139)
```

`-A login` – add a rule to the chain called `login`

`-s` – specifies the source network and subnet mask

`-d` – specifies the destination network and subnet mask

`-p` – specifies the protocol to be used (1 = ICMP, 6 = TCP, 17 = UDP)

`-j` – jumps to the appropriate chain or rule (Deny, accept, reject)

`-l` – enables logging on this rule

RPC and NFS-- Portmap/rpcbind (111/tcp and 111/udp), NFS (2049/tcp and 2049/udp), lockd (4045/tcp and

4045/udp)

RPC

as defined at "whatis.com":

RPC (Remote Procedure Call) is a protocol that one program can use to request a service from a program located in another computer in a network without having to understand network details. (A *procedure call* is also sometimes known as a *function call* or a *subroutine call*.) RPC uses the client/server model. The requesting program is a client and the service-providing program is the server. Like a regular or local procedure call, an RPC is a synchronous operation requiring the requesting program to be suspended until the results of the remote procedure are returned. However, the use of *lightweight processes* or threads that share the same address space allows multiple RPCs to be performed concurrently.

When program statements that use RPC are compiled into an executable program, a stub is included in the compiled code that acts as the representative of the remote procedure code. When the program is run and the procedure call is issued, the stub receives the request and forwards it to a client runtime program in the local computer. The client runtime program has the knowledge of how to address the remote computer and server application and sends the message across the network that requests the remote procedure. Similarly, the server includes a runtime program and stub that interface with the remote procedure itself. Results are returned the same way.

There are several RPC models and implementations. A popular model and implementation is the Open Software Foundations' Distributed Computing Environment (DCE). The IEEE defines RPC in its *ISO Remote Procedure Call Specification*, ISO/IEC CD 11578 N6561, ISO/IEC, November 1991.

RPC spans the transport layer and the application layer in the Open Systems Interconnection (OSI) model of network communication. RPC makes it easier to develop an application that includes multiple programs distributed in a network.

Possible vulnerabilities associated with RPC

There are many different vulnerabilities associated with RPC. Because RPCs allow a user or program to execute an application or code on a remote computer there are serious consequences that can arise.

Malicious code can be inserted into an RPC to provide root access or create user accounts.

Consequences associated with blocking RPC

Many databases and web servers communicate with each other via RPC. If you have multiple MS-Exchange servers or other MS servers they will use RPCs to communicate different types of data and execute parts of an application.

NFS

as defined at "whatis.com":

The Network File System (NFS) is a client/server application that lets a computer user view and optionally store and update files on a remote computer as though they were on the user's own computer. The user's system needs to have an NFS client and the other computer needs the NFS server. Both of them require that you also have TCP/IP installed since the NFS server and client use TCP/IP as the program that sends the files and updates back and forth. (However, the User Datagram Protocol, UDP, which comes with TCP/IP, is used instead of TCP with earlier versions of NFS.)

NFS was developed by Sun Microsystems and has been designated a file server standard. Its protocol uses the Remote Procedure Call (RPC) method of communication between computers. You can install NFS on Windows 95 and some other operating systems using products like Sun's Solstice Network Client.

Using NFS, the user or a system administrator can *mount* all or a portion of a *file system* (which is a portion of the hierarchical tree in any file directory and subdirectory, including the one you find on your PC or Mac). The portion of your file system that is mounted (designated as accessible) can be accessed with whatever privileges go with your access to each file (read-only or read-write).

Possible vulnerabilities associated with NFS

NFS provides file sharing across a network. It allows a user to mount drives across the network and use them as if they were locally installed in the machine. Attackers can exploit the trust relations between client and server to view, or manipulate files on a remote machine.

Consequences associated with blocking NFS

By eliminating NFS you are disabling the ability to remotely mount files across the firewall. This could impact moving large files to an FTP or web server.

Lockd

as defined in the man pages:

The lockd utility is part of the NFS lock manager, which supports record locking operations on NFS files. See `fcntl(2)` and `lockf(3C)`. The lock manager provides two functions: o it forwards `fcntl(2)` locking requests for NFS

mounted file systems to the lock manager on the NFS server or it generates local file locking operations in response to requests forwarded from lock managers running on NFS client machines. State information kept by the lock manager about these locking requests can be lost if the lockd is killed or the operating system is rebooted. Some of this information can be recovered as follows. When the server lock manager restarts, it waits for a grace period for all client-site lock managers to submit reclaim requests. Client-site lock managers, on the other hand, are notified by the status monitor daemon, statd(1M), of the restart and promptly resubmit previously granted lock requests. If the lock daemon fails to secure a previously granted lock at the server site, then it sends SIGLOST to a process.

Possible vulnerabilities associated with lockd

There have been several exploits that allow the remote execution of commands by using exploits found in the sun implementation of this. Some patches are available but are not sufficient in some cases. The linux version also has known vulnerabilities with this service that can lead to a DoS attack.

Consequences associated with blocking lockd

This service is unneeded unless you are allowing NFS through the firewall. Because we are not allowing NFS through this service is unnecessary.

```
-A rpc -s 0/0 -d 95.5.6.0/24 111 -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 2049 -y -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 4045 -y -p 6 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 111 -p 17 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 2049 -p 17 -j DENY -l
-A rpc -s 0/0 -d 95.5.6.0/24 4045 -p 17 -j DENY -l
```

NetBIOS in Windows NT -- 135 (tcp and udp), 137 (udp), 138 (udp), 139 (tcp). Windows 2000 - earlier ports plus 445(tcp and udp)

NetBIOS in Windows NT and Windows 2000

Possible vulnerabilities associated with NetBIOS in Windows NT and Windows 2000

Because of the networking stack implemented in the windows NT/2000 operating systems vulnerabilities associated with the NetBIOS implementation can lead to denial of Service attacks as interception of data.

Consequences associated with blocking NetBIOS in Windows NT and Windows 2000

NetBIOS communications between the Internet and internal networks will not function. This could limit the types of servers and implementations of services available through the firewall to the Internet.

```
-A netbios -s 0/0 -d 95.5.6.0/24 135 -p 6 -j DENY -l
-A netbios -s 0/0 -d 95.5.6.0/24 135 -p 17 -j DENY -l
-A netbios -s 0/0 -d 95.5.6.0/24 137:139 -p 6 -j DENY -l
-A netbios -s 0/0 -d 95.5.6.0/24 137:139 -p 17 -j DENY -l
```

X Windows -- 6000/tcp through 6255/tcp

X Windows

as defined at "whatis.com":

The X Window System (sometimes referred to as "X" or as "XWindows") is an open, cross-platform, client-server system for managing a windowed graphical user interface in a distributed network. In general, such systems are known as windowing systems. In X Window, the client-server relationship is reversed from the usual. Remote computers contain applications that make client requests for display management services in each PC or workstation. X Window is primarily used in networks of interconnected mainframe, minicomputer, and workstations. It is also used on the X terminal, which is essentially a workstation with display management capabilities but without its own applications. (The X terminal can be seen as a predecessor of the network PCs or thin client computers that are now beginning to be widely installed.)

The X Window System was the result of research efforts in the early 1980s at Stanford University and MIT, aided by IBM, to develop a platform-independent graphics protocol. The X Window System is an open standard that is managed by the X.Org consortium. Although Microsoft has its own platform-dependent windowing system (an

integral part of the Windows 95/98/NT operating systems), there are vendor-supplied X Windows products that can be installed to run on these systems.

Possible vulnerabilities associated with X Windows

It is possible that X servers are improperly configured which can allow any user to connect to its services. This could result in root level compromises because administrators use these services to administrate their systems. Some administrators configure their SSH clients to forward X connections through their encrypted tunnel. This is not the default and some sites do not have this available to them.

Consequences associated with blocking X Windows

When blocking these services there is potential to also block legitimate connections because they use ephemeral ports (ports above 1023). Clients use these ports when making connections to servers. If you do not specify any flags in the TCP header then all connections with a destination port between 6000 and 6255 will be blocked, including responses to valid internal user requests. When using ipchains one can specify the `-y` option to specify the syn flag being set. When the syn flag is set that indicates a connection attempt that means this is not a response to a user request.

```
-A XWin -s 0/0 -d 95.5.6.0/24 6000:6255 -y -p 6 -j DENY -l
```

Naming services-- DNS (53/udp) to all machines which are not DNS servers, DNS zone transfers (53/tcp) except from external secondaries, LDAP (389/tcp and 389/udp)

DNS

as defined at "whatis.com":

The domain name system (DNS) is the way that Internet domain names are located and translated into IP (Internet Protocol) addresses. A domain name is a meaningful and easy-to-remember "handle" for an Internet address. Because maintaining a central list of domain name/IP address correspondences would be impractical, the lists of domain names and IP addresses are distributed throughout the Internet in a hierarchy of authority. There is probably a DNS server within close geographic proximity to your access provider that maps the domain names in your Internet requests or forwards them to other servers in the Internet.

Possible vulnerabilities associated with DNS

DNS Zone Transfer

A DNS Zone Transfer occurs when a DNS server transfers its entire zone database to another DNS server. These transfers should occur between trusted servers only (primary and secondary DNS server). These zones can contain all or some records associated with a corporation's infrastructure. These can include e-mail, web, file and print servers.

Possible vulnerabilities associated with DNS and DNS Zone Transfers:

Most DNS servers are powered by a program called BIND. This program is very old and complex which usually leads to a number of vulnerabilities such as root compromises and buffer overruns. DNS Zone Transfers provide attackers information such as computer name and IP address translations. This could allow an attacker information to refine their attack without using certain mapping or probing programs which may trigger an alert on many intrusion detection systems.

Consequences associated with blocking DNS and DNS Zone transfers:

When blocking these services you are attempting to minimize your exposure to these attacks. However, by creating these filters you may need to implement both internal and external only DNS servers. This could lead to additional computers which need updating and software patches.

LDAP

as defined at "whatis.com":

LDAP (Lightweight Directory Access Protocol) is a software protocol for enabling anyone to locate organizations, individuals, and other resources such as files and devices in a network, whether on the Internet or on a corporate intranet. LDAP is a "lightweight" (smaller amount of code) version of DAP (Directory Access Protocol), which is part of X.500, a standard for directory services in a network. LDAP is lighter because in its initial version it did not include security features. LDAP originated at the University of Michigan and has been endorsed by at least 40 companies. Netscape includes it in its latest Communicator suite of products. Microsoft includes it as part of what it

calls Active Directory in a number of products including Outlook Express. Novell's NetWare Directory Services interoperates with LDAP. Cisco also supports it in its networking products.

In a network, a directory tells you where in the network something is located. On TCP/IP networks (including the Internet), the Domain Name System (DNS) is the directory system used to relate the domain name to a specific network address (a unique location on the network). However, you may not know the domain name. LDAP allows you to search for an individual without knowing where they're located (although additional information will help with the search).

An LDAP directory is organized in a simple "tree" hierarchy consisting of the following levels:

- The "root" directory (the starting place or the source of the tree), which branches out to
- Countries, each of which branches out to
- Organizations, which branch out to
- Organizational units (divisions, departments, and so forth), which branches out to (includes an entry for)
- Individuals (which includes people, files, and shared resources such as printers)

An LDAP directory can be distributed among many servers. Each server can have a replicated version of the total directory that is synchronized periodically. An LDAP server is called a Directory System Agent (DSA). An LDAP server that receives a request from a user takes responsibility for the request, passing it to other DSAs as necessary, but ensuring a single coordinated response for the user.

Possible vulnerabilities associated with LDAP

There are a number of servers that have security holes that can lead to buffer overflow conditions due to their LDAP implementations. LDAP services can also provide an attacker with a great deal of useful information if it is not properly secured. If an attacker were to compromise an LDAP server it is possible for that attacker to gain intimate knowledge of an organization's structure, trust relations, and user base.

Consequences associated with blocking LDAP

Servers and services that require LDAP to function will not work. If the corporation has implemented a worldwide LDAP system VPNs or other types of private connections will be needed to continue to function.

```
-A name -s 0/0 -d 95.5.6.0/24 389 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 389 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.10/32 53 -p 17 -j ACCEPT
```

Mail-- SMTP (25/tcp) to all machines, which are not external mail relays, POP (109/tcp and 110/tcp), IMAP (143/tcp)

SMTP

as defined at "whatis.com":

SMTP (Simple Mail Transfer Protocol) is a TCP/IP protocol used in sending and receiving e-mail. However, since it's limited in its ability to queue messages at the receiving end, it's usually used with one of two other protocols, POP3 or IMAP, that let the user save messages in a server mailbox and download them periodically from the server. In other words, users typically use a program that uses SMTP for sending e-mail and either POP3 or IMAP for receiving messages that have been received for them at their local server. Most mail programs such as Eudora let you specify both an SMTP server and a POP server. On UNIX-based systems, sendmail is the most widely-used SMTP server for e-mail. A commercial package, Sendmail, includes a POP3 server and also comes in a version for Windows NT.

SMTP usually is implemented to operate over TCP port 25. The details of SMTP are in RFC 821 of the Internet Engineering Task Force (IETF). An alternative to SMTP that is widely used in Europe is X.400.

Possible vulnerabilities associated with SMTP

Unsecured SMTP servers can allow open relaying. Sendmail is a program that is widely used on the Internet for e-mail. It is very old and complicated. There are a number of vulnerabilities that have been identified in sendmail of all versions. Knowing what version of the software is running can provide an attacker much information that can lead to buffer overflows and other attacks that can lead to a root compromise. While you can not control who knows what version of the software you are running

(simply telnet to port 25 on the target) you can control which servers are available to the internet which is what this rule attempts to do.

Consequences associated with blocking SMTP

This could complicate implementation of an enterprise and Internet e-mail system. However, because a corporation only has one bridgehead server this usually isn't difficult to plan for.

POP3

as defined at "whatis.com":

POP3 (Post Office Protocol 3) is the most recent version of a standard protocol for receiving e-mail. POP3 is a client/server protocol in which e-mail is received and held for you by your Internet server. Periodically, you (or your client e-mail receiver) check your mailbox on the server and download any mail. POP3 is built into the Netmanager suite of Internet products and one of the most popular e-mail products, Eudora. It's also built into the Netscape and Microsoft Internet Explorer browsers.

Possible vulnerabilities associated with POP3

There are a number of known buffer overflow and remote access vulnerabilities with POP3 and IMAP daemons and services. These vulnerabilities can lead to DoS attacks as well as system root compromises. Because these are client software packages it is also possible that email can be read if users have chosen poor passwords.

Consequences associated with blocking POP3

POP3 is a client/server application. It is unneeded unless employees must access their e-mail from the Internet. If access from the Internet to user accounts is necessary then protocol bay also be necessary

IMAP

as defined at "whatis.com":

IMAP (Internet Message Access Protocol) is a standard protocol for accessing e-mail from your local server. IMAP (the latest version is IMAP4) is a client/server protocol in which e-mail is received and held for you by your Internet server. You (or your e-mail client) can view just the heading and the sender of the letter and then decide whether to download the mail. You can also create and manipulate folders or mailboxes on the server, delete messages, or search for certain parts or an entire note. IMAP requires continual access to the server during the time that you are working with your mail.

IMAP can be thought of as a remote file server. POP can be thought of as a "store-and-forward" service.

POP and IMAP deal with the receiving of e-mail from your local server and are not to be confused with SMTP, a protocol for transferring e-mail between points on the Internet. You send e-mail with SMTP and a mail handler receives it on your recipient's behalf. Then the mail is read using POP or IMAP.

Possible vulnerabilities associated with IMAP

There are a number of known buffer overflow and remote access vulnerabilities with POP3 and IMAP daemons and services. These vulnerabilities can lead to DoS attacks as well as system root compromises. Because these are client software packages it is also possible that email can be read if users have chosen poor passwords.

Consequences associated with blocking IMAP

IMAP is a client/server application. It is unneeded unless employees must access their e-mail from the Internet. If access from the Internet to user accounts is necessary then protocol bay also be necessary

```
-A mail -s 0/0 1024:65535 -d 95.5.6.25/32 25 -p 6 -j ACCEPT
-A mail -s 95.5.6.25/32 1024:65535 -d 0/0 25 -p 6 -j ACCEPT
-A mail -s 95.5.6.25/32 1024:65535 -d 172.16.1.25/32 25 -p 6 -j ACCEPT
-A mail -s 172.16.1.25/32 1024:65535 -d 95.5.6.25/32 25 -p 6 -j ACCEPT
-A mail -s 0/0 -d 95.5.6.0/24 25 -p 6 -j DENY -l
-A mail -s 0/0 -d 95.5.6.0/24 109:110 -p 6 -j DENY -l
-A mail -s 0/0 -d 95.5.6.0/24 143 -p 6 -j DENY -l
```

Web-- HTTP (80/tcp) and SSL (443/tcp) except to external Web servers, may also want to block common high-order HTTP port choices (8000/tcp, 8080/tcp, 8888/tcp, etc.)

HTTP

as defined at "whatis.com":

The Hypertext Transfer Protocol (HTTP) is the set of rules for exchanging files (text, graphic images, sound, video, and other multimedia files) on the World Wide Web. Relative to the TCP/IP suite of protocols (which are the basis for information exchange on the Internet), HTTP is an application protocol.

Essential concepts that are part of HTTP include (as its name implies) the idea that files can contain references to other files whose selection will elicit additional transfer requests. Any Web server machine contains, in addition to the HTML and other files it can serve, an HTTP daemon, a program that is designed to wait for HTTP requests and handle them when they arrive. Your Web browser is an HTTP client, sending requests to server machines. When the browser user enters file requests by either "opening" a Web file (typing in a Uniform Resource Locator or URL) or clicking on a hypertext link, the browser builds an HTTP request and sends it to the Internet Protocol address indicated by the URL. The HTTP daemon in the destination server machine receives the request and, after any necessary processing, the requested file is returned.

The latest version of HTTP is HTTP 1.1.

Possible vulnerabilities associated with HTTP

When web programmers create applications for use via http they rarely perform extensive error checking and therefore are susceptible to buffer overruns and other exploits. When creating a secure environment one should attempt to minimize exposure to outside entities such as the Internet. By restricting what computers and services an attacker can reach from the Internet one can reduce the

Consequences associated with blocking (restricting) HTTP

Because it is very easy to create a web site and web server it is advisable to limit what servers are accessible from the Internet. Rogue web servers will not function outside your network when this protocol is being restricted.

SSL

as defined at "whatis.com":

SSL (Secure Sockets Layer) is a program layer created by Netscape for managing the security of message transmissions in a network. Netscape's idea is that the programming for keeping your messages confidential ought to be contained in a program layer between an application (such as your Web browser or HTTP) and the Internet's TCP/IP layers. The "sockets" part of the term refers to the sockets method of passing data back and forth between a client and a server program in a network or between program layers in the same computer. Netscape's SSL uses the public-and-private key encryption system from RSA, which also includes the use of a digital certificate.

SSL is an integral part of each Netscape browser. If a Web site is on a Netscape server, SSL can be enabled and specific Web pages can be identified as requiring SSL access. Other servers can be enabled by using Netscape's SSLRef program library, which can be downloaded for noncommercial use or licensed for commercial use.

Netscape has offered SSL as a proposed standard protocol to the World Wide Web Consortium (W3C) and the Internet Engineering Task Force (IETF) as a standard security approach for Web browsers and servers.

Possible vulnerabilities associated with SSL

Because of the relationship between ssl and http the vulnerabilities are very similar. It is also possible that if ssl is not configured properly it can provide a false sense of security. Not all transactions may be encrypted during a given session.

Consequences associated with blocking (restricting) SSL

Rogue web servers will not function outside your network when this protocol is being restricted.

```
-A web -s 0/0 -d 95.5.6.20/32 80 -p 6 -j ACCEPT
-A web -s 0/0 -d 95.5.6.20/32 443 -p 6 -j ACCEPT
-A web -s 0/0 -d 95.5.6.0/24 80 -p 6 -j DENY -1
-A web -s 0/0 -d 95.5.6.0/24 8080 -y -p 6 -j DENY -1
-A web -s 0/0 -d 95.5.6.0/24 8888 -y -p 6 -j DENY -1
```

"Small Services"-- ports below 20/tcp and 20/udp, time (37/tcp and 37/udp)

These services were originally used for troubleshooting connections and server functionality when networks were in their infancy. These services include echo (which responds back with the input characters), chargen (character generator), and daytime (returns the date and time on the remote system). There are many others

I've chosen just to highlight a few. These services are very susceptible to DoS attacks due to their nature. An attacker only needs to send a charge packet which initiates a loop with echo and this will put the traffic in a loop where charge creates a character and sends it to echo which then repeats that character.

```
-A small -s 0/0 0/0 0:20 -p 6 -j DENY -l
-A small -s 0/0 0/0 0:20 -p 17 -j DENY -l
-A small -s 0/0 0/0 37 -p 6 -j DENY -l
-A small -s 0/0 0/0 37 -p 17 -j DENY -l
```

Miscellaneous-- TFTP (69/udp), finger (79/tcp), NNTP (119/tcp), NTP (123/tcp), LPD (515/tcp), syslog (514/udp), SNMP (161/tcp and 161/udp, 162/tcp and 162/udp), BGP (179/tcp), SOCKS (1080/tcp)

TFTP

as defined at "whatis.com":

TFTP (Trivial File Transfer Protocol) is a network application that is simpler than the File Transfer Protocol (FTP) but less capable. It is used where user authentication and directory visibility are not required. TFTP uses the User Datagram Protocol (UDP) rather than the Transmission Control Protocol (TCP). TFTP is described formally in RFC 1350.

Possible vulnerabilities associated with TFTP

TFTP has no authentication associated with it. Generally network hardware devices such as routers, switches, and network printers use this service to backup configuration files. If an attacker were to change that configuration file they could cause a denial of service or other type of attack. An attacker could also read the configuration file and gain a better understanding of the network architecture and identify systems or services to attack.

Consequences associated with blocking TFTP

Making backups of external routers, or switches may require a machine to temporarily be configured with tftp so that they can send the information to it. This could be a laptop configured so that it may be placed outside the firewall for the duration of the backup (usually 2-3 minutes only).

Finger

As defined at "whatis.com":

Finger is a program that tells you the name associated with an e-mail address. It may also tell you whether they are currently logged on at their system or their most recent logon session and possibly other information, depending on the data that is maintained about users on that computer. Finger originated as part of BSD UNIX.

To finger another Internet user, you need to have the finger program on your computer or you can go to a finger gateway on the Web and enter the e-mail address. The server at the other end must be set up to handle finger requests. A ".plan" file can be created for any user that can be fingered. Commonly, colleges, universities, and large corporations set up a finger facility. Your own Internet access provider may also set up information about you and other subscribers that someone else can "finger." (To find out, enter your own e-mail address at a finger gateway.)

Possible vulnerabilities associated with finger

Finger provides any user information that can be dangerous to the company and the users. Finger will provide information such as most recent logon and length of the session. This will allow a hacker to identify a user's habits such as arrival and departure time at work. While this information may not be dangerous by itself, it can provide a timetable to base an attack on the person or their computer.

Consequences associated with blocking Finger

This information will not be available from the Internet. This should not be a concern since no services or applications depend on this to function properly.

NNTP

as defined at "whatis.com":

NNTP (Network News Transfer Protocol) is the predominant protocol used by computers (servers and clients) for managing the notes posted on Usenet newsgroups. NNTP replaced the original Usenet protocol, UNIX-to-UNIX Copy Protocol (UUCP) some time ago. NNTP servers manage the global network of collected Usenet newsgroups and include the server at your Internet access provider. An NNTP client is included as part of a Netscape, Internet Explorer, Opera, or other Web browser or you may use a separate client program called a newsreader.

Possible vulnerabilities associated with NNTP

Several NNTP servers have vulnerabilities that can lead to buffer overflow and DoS attacks. Also the amount of traffic and information that can be passed to a NNTP server could fill a drive array thus causing a DoS attack. Because these systems are now being integrated into e-mail packages all of the vulnerabilities associated with e-mail clients can also be exploited here as well.

Consequences associated with blocking NNTP

This rule would have to be modified if this network was to begin hosting its own externally available news server. An internal only NNTP server is not affected by this rule.

NTP

as defined at "whatis.com":

Network Time Protocol (NTP) is a protocol that is used to synchronize computer clock times in a network of computers. Developed by David Mills at the University of Delaware, NTP is now an Internet standard. In common with similar protocols, NTP uses Coordinated Universal Time (UTC) to synchronize computer clock times to a millisecond, and sometimes to a fraction of a millisecond.

Accurate time across a network is important for many reasons; even small fractions of a second can cause problems. For example, distributed procedures depend on coordinated times to ensure that proper sequences are followed. Security mechanisms depend on coordinated times across the network. File system updates carried out by a number of computers also depend on synchronized clock times. Air traffic control systems provide a graphic illustration of the need for coordinated times, since flight paths require very precise timing (imagine the situation if air traffic controller computer clock times were not synchronized).

UTC time is obtained using several different methods, including radio and satellite systems. Specialized receivers are available for high-level services such as the Global Positioning System (GPS) and the governments of some nations. However, it is not practical or cost-effective to equip every computer with one of these receivers. Instead, computers designated as *primary time servers* are outfitted with the receivers and they use protocols such as NTP to synchronize the clock times of networked computers. Degrees of separation from the UTC source are defined as strata. A radio clock (which receives true time from a dedicated transmitter or satellite navigation system) is stratum-0; a computer that is directly linked to the radio clock is stratum-1; a computer that receives its time from a stratum-1 computer is stratum-2, and so on.

The term NTP applies to both the protocol and the client/server programs that run on computers. The programs are compiled by the user as an NTP client, NTP server, or both. In basic terms, the NTP client initiates a time request exchange with the time server. As a result of this exchange, the client is able to calculate the link delay, its local offset, and adjust its local clock to match the clock at the server's computer. As a rule, six exchanges over a period of about five to 10 minutes are required to initially set the clock. Once synchronized, the client updates the clock about once every 10 minutes, usually requiring only a single message exchange. Redundant servers and varied network paths are used to ensure reliability and accuracy. In addition to client/server synchronization, NTP also supports broadcast synchronization of peer computer clocks. NTP is designed to be highly fault-tolerant and scalable.

Possible vulnerabilities associated with NTP

If someone can spoof a timeserver (usually no authentication or verification with NTP) then it is possible to change times on a specific system. This can have an impact on systems that rely on timestamps of certificates or passwords. For example, Windows 2000 relies on a time stamp associated with a changed password to ensure that the password database doesn't corrupt. If the time on a domain controller were to be changed then the domain password files could be corrupted and cause problems with people authenticating. Replication between database servers could also be affected if the time were to become corrupt.

Consequences associated with blocking NTP

An internal timeserver may be required so time can be synchronized within the organization. This should be a system that can synchronize with a known good time source such as the NIST time sources across the country.

LPD

LPD is the daemon used for the line printer protocol used by most versions of UNIX/Linux for printing. It listens on TCP port 514 and only accepts connections from TCP ports 721-731.

Possible vulnerabilities associated with LDP

Certain exploits can take advantage of this service allowing the attacker to run arbitrary code in the

context of the user lp. Lp is the account associated with this daemon on some versions of UNIX/Linux.

Consequences associated with blocking LDP

This service is used for printing only and should be confined to the internal network only. By blocking this service at the firewall, servers and clients will not be able to print to/from devices on the other side of the firewall.

Syslog

Syslogd is a program that is used extensively for system logging on UNIX/Linux servers. Many UNIX programs can utilize the syslog capabilities for reporting and logging. Syslog can communicate between servers to create a centralized logging service. Syslog can create entries based on several different levels of criteria such as info, error, alert, warning, and critical.

Possible vulnerabilities associated with syslog

There are several buffer overflow attacks that target the syslogd. If one of the centralized servers were discovered it is possible to attack that server and take it off line which could severely limit logging capabilities on some networks which rely on a single system. If the centralized server were to go off line an attacker could go undetected while installing trojans or altering system software on other network systems.

Consequences associated with blocking syslog

Logs from servers outside can not be collected through the firewall. There are other ways for collecting those logs such as via a serial connection to another server. This would provide a level of security because an attacker would have to compromise the external machine before they would be able to attack the internal syslog server.

SNMP

as defined at "whatis.com":

SNMP is the protocol governing network management and the monitoring of network devices and their functions. It is not necessarily limited to TCP/IP networks.

Possible vulnerabilities associated with SNMP

The main form of authentication in SNMP is the community name. Many systems are installed using the default community name "public" some administrators change the name to "private". If a hacker were able to use SNMP on a machine they could do serious damage to a computer, or network. They could change the IP address of a router, the route or default gateway associated with a specific route or the server name on a local server.

Consequences associated with blocking SNMP

Management of certain systems using SNMP will not be available through the Internet. This is not a significant consequence because of the limited number of systems outside the firewall.

BGP

as defined at "whatis.com":

BGP (Border Gateway Protocol) is a protocol for exchanging routing information between gateway hosts (each with its own router) in a network of autonomous systems. BGP is often the protocol used between gateway hosts on the Internet. The routing table contains a list of known routers, the addresses they can reach, and a cost metric associated with the path to each router so that the best available route is chosen.

Hosts using BGP communicate using the Transmission Control Protocol (TCP) and send updated router table information only when one host has detected a change. Only the affected part of the routing table is sent. BGP-4, the latest version, lets administrators configure cost metrics based on policy statements. (BGP-4 is sometimes called BGP4, without the hyphen.)

BGP communicates with autonomous (local) networks using Internal BGP (IBGP) since it doesn't work well with IGP. The routers inside the autonomous network thus maintain two routing tables: one for the interior gateway protocol and one for IBGP.

BGP-4 makes it easy to use Classless Inter-Domain Routing (CIDR), which is a way to have more addresses within the network than with the current Internet Protocol address assignment scheme.

Possible vulnerabilities associated with BGP

BGP does not use a strong identification and authentication mechanism. Because of this there is the

possibility that a router could poison the routing table of a router. Because this is an external gateway protocol there is no reason it should be traveling inside a corporate network. If this port is open an attacker could use this port as a hole in which they send trojan traffic or other malicious code through.

Consequences associated with blocking BGP

Because BGP is used for communicating router updates at gateway devices it should not be necessary on an internal network. Network engineers should be aware of this and configure the internal routers appropriately so that the routers do not attempt to use this protocol.

SOCKS

as defined at "whatis.com":

Socks (or "SOCKS") is a protocol that a proxy server can use to accept requests from client users in a company's network so that it can forward them across the Internet. Socks uses sockets to represent and keep track of individual connections. The client side of Socks is built into certain Web browsers and the server side can be added to a proxy server.

A socks server handles requests from clients inside a company's firewall and either allows or rejects connection requests, based on the requested Internet destination or user identification. Once a connection and a subsequent "bind" request have been set up, the flow of information exchange follows the usual protocol (for example, the Web's HTTP protocol).

Possible vulnerabilities associated with SOCKS

There have been a few DoS attacks regarding SOCKS proxies that will fail if a request is sent with especially long arguments.

Consequences associated with blocking SOCKS

There is only an issue if SOCKS is required on the network. This is not normally the case and it is certainly possible to use a regular web, or ftp proxy instead of attempting to use SOCKS

```
-A name -s 0/0 -d 95.5.6.0/24 79 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 69 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 119 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 123 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 515 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 161:162 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 161:162 -p 17 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 179 -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 1080 -y -p 6 -j DENY -l
-A name -s 0/0 -d 95.5.6.0/24 514 -p 17 -j DENY -l
```

ICMP-- block incoming echo request (ping and Windows traceroute), block outgoing echo replies, time exceeded, and unreachable messages

ICMP

as defined at "whatis.com":

ICMP is a message control and error-reporting protocol between a host server and a gateway to the Internet. ICMP uses Internet Protocol (IP) datagrams, but the messages are processed by the IP software and are not directly apparent to the application user.

Possible vulnerabilities associated with ICMP

There are several attacks that can use ICMP. One of the more notables is the smurf or smurfamp attack that is a DoS or DDoS attack. ICMP also provides information back to systems such as redirect information. ICMP can also provide a useful tool when mapping a network. One can use traceroute or tracert to map the routers a packet takes to reach a host or network. Once a particular router or gateway has been identified a DoS attack can be launched against it.

Consequences associated with blocking ICMP

The services we are blocking will disable incoming ping traffic (people will not be able to ping any internal hosts because the request is denied). Outgoing responses are also blocked so that other DoS attacks will also be disabled. The responses will mitigate the possibility that internal machines are used

as part of a DDoS attack.

```
-A icmp -s 0/0 -d 95.5.6.0/24 --icmp-type echo-request -p 1 -j DENY
-A icmp -s 95.5.6.0/24 -d 0/0 --icmp-type echo-reply -p 1 -j DENY
-A icmp -s 0/0 -d 0/0 --icmp-type time-exceeded -p 1 -j DENY
-A icmp -s 0/0 -d 0/0 --icmp-type destination-unreachable -p 1 -j DENY
```

Testing these filters and rules:

These rules can be tested in a number of different ways. The easiest way would be to first run a program like Nessus or CyberCop from outside the network (and border router) against the IP address range of 95.5.6.0/24. This will help determine whether some of these rules are not successful in filtering the proper traffic. What these automated tools will tell you is whether the rule has failed or if it is successful. These tools use a database of tests that they use to determine whether a vulnerability exists. If one matches the database it will alert the user. If a match does not exist it simply means that there was no match in the database, not that a vulnerability does not exist. If both programs were to be used it would help identify certain filters that may be of concern. I would also recommend attempting to exploit one or more of these services by using the documented exploits available on such sites as www.securityfocus.com, www.rootshell.com, and packetstorm.securify.com. This is a time-consuming process that can not be overlooked. It should also be an ongoing process that occurs on a regular basis when exploits are discovered or services are added.

Another test could be to use a program like hping2 to craft packets which are designed to avoid being filtered or detected. Execute this program from an external location such as a dialup account. You should have a sniffer or IDS system setup to look for the traffic getting through to e0 (internal interface). If the sniffer or IDS picks up the traffic then the filters are not working properly.

© SANS Institute 2000 - 2002