Global Information Assurance Certification Paper

Quei-Len Lee

GIAC Firewall and Perimeter Protection Curriculum
Practical Assignment for SANS Security DC 2000
July 5 – 10, 2000
Version 1.1

Write a tutorial on how to implement each recommended action in the filtering policy below on your firewall or perimeter defense solution.

Assignment 1 – 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.

Introduction

Spoofing is a technique where the IP address of a packet is altered to make it appear as though the packet has originated from your internal network. The target system can be fooled into believing that the attacker's machine is really a trusted machine. IP source route, on the other hand, allows someone to specify an explicit path to the destination, overriding the usual route.

IP spoofing and source route can be used for malicious attacks. Here is a scenario of such an attack. A spoofing host can send a packet, with bogus source IP address, pretending to be a trusted host to a target host in a network that allows source routing. Normally, if a network does not allow source routing, any response from the target host will be sent back using the normal route, i.e. to the real trusted host. When a network allows source routing, the target host will send responses back, based on the path dictated by the spoofed packet. In that case, protocols that depend on the IP source address for authentication can be exploited to penetrate the target system. For example, many UNIX 'r' commands can be configured to allow access to trusted hosts without a password. The outcome is that a spoofing host has gained unauthorized access to the trusted hosts.

The simplest method to defend against the source routing problem is to simply reject packets with the option. On Cisco IOS, the command is "no ip source-route". It is extremely important that you verify that the source route options are disabled on your router.

IP spoofing can be diminished through the configurations of your router by blocking all packets coming from outside with the following source IP addresses:

your internal IP,
private IP address space (10.*.*.*, 172.16.*.* - 172.31.255.255, 192.168.*.*)
and the 127 addresses.

The filtering of inbound network traffic for spoofed packets is called an 'anti-spoofing' or 'ingress' filter. An example of an 'ingress' router command is shown in the section - 'Syntax of filter'.

Most firewall systems also include anti-spoofing features. For CheckPoint Firewall-1, you can define anti-spoofing in the property window of the network object. For example, if a gateway is to enforce anti-spoofing, the spoof tracking parameters are defined in the gateway's workstation properties window.

It is also important to note that the filtering of outbound packets can ensure that your router and/or firewall forward only IP packets with correct source IP address for your network. While the egress filters do not protect your network from spoofing, it protects you from having your network used as a DDOS (Distributed Denial of Service) source and damaging other networks. The filtering of packets outbound from a private network is called 'egress' filter. An example of 'egress' filter can be found below.

Syntax of the filter

For Cisco's IOS, the command to disable IP source route is:

```
no ip source-route
```

Given a Cisco router that routes for a class-B address space 128.204.*.*, the 'anti-spoofing' or 'ingress' filter is the following.

```
! Interface serial 0 connects to the Internet
Interface serial 0
        ip address 128.204.77.1  255.255.0.0
        ip address-group 11 in

! Deny private addresses
access-list 11 deny 192.168.0.0 0.0.255.255
access-list 11 deny 172.16.0.0 0.15.255.255
access-list 11 deny 10.0.0.0 0.255.255.255
! Deny IP with localhost – 127.0.0.1
access-list 11 deny 127.0.0.0 0.255.255.255
! Deny incoming packets with internal addresses and log the attempts
access-list 11 deny 128.204.0.0 0.0.255.255 log
```

Given a Cisco router that routes for a class-C private address space 192.168.1.*, the 'egress' ACL is the following:

```
Interface ethernet 0
        ip address 192.168.1.1 255.255.255.0
        ip access-group 11 in

! Allow only outbound packets that have internal addresses
```

access-list 11 permit 192.168.1.0 0.0.0.255
access-list 11 deny any log

Description of each of the parts of the filter

"Anti-spoofing" or "ingress" filter -
In the example of "anti-spoofing" or "ingress" filter, access-list 11 is a Cisco "standard" access-list, i.e. the access-list numbers with a value between 1-99. It indicates to the router that only the source IP address of a packet needs to be examined. The access-list is to be applied to the router on inbound packets to the interface serial 0. The first 3 access-list filters deny any packets with source IP in the private address space (i.e. 192.168.*.*, 172.16.*.*-172.31.255.255 and 10.*.*.*). The 4th access-list denies packets with localhost (127.0.0.1) . The 5th access-list indicates to the router that any inbound address using my internal network space (128.204.*.*), should be denied and logged.

"Egress" filter -
In the example of "egress" filter, the access-list indicates to the router, all incoming packets to the interface ethernet 0 should have the source IP of my internal network - 192.168.1.*. Any packet leaving my network that has a source IP does not belong to the internal network, should be dropped and logged.

Explain how to apply the filter

The no source-route command and ingress and egress filters are applied to the Cisco router by entering the commands as shown in the examples above at the router's command line interface.

Explain how to test the filter

The testing of the ingress filter for packets coming from outside with private addresses or internal addresses (in this example, a source IP address with the prefix 216.80.17) can be done using one of the freely available tools on the Internet. It's simply a matter of generating a packet with an internal address and sending it from outside. The "ingress" filter indicates to the router that any packets coming from outside, that have our internal addresses, should be denied and logged. The denied packets will be found in the logs.

The testing of the egress filter can be done by setting up a host with a non-internal IP address on the internal network and then generate some packets destined to the Internet. You can also use one of the freely available tools to craft a packet with a spoofed source IP address. Since the egress filter tells the router to log those denied packets, you will be able to find the denied outbound attempts in the log file.

Assignment 2 – Login Services – telnet (23/tcp), SSH (22/tcp), FTP (21/tcp), NetBIOS (139/tcp), rlogin et al (512/tcp through 514/tcp).

Introduction

Telnet and rlogin are applications that provide simple terminal access to a machine. Both telnet and rlogin authenticate via a login name and password. Telnet sessions often come from an 'untrusted' system. Rlogin can be configured to allow access from a remote host without a password, if the remote machine is included in a 'trusting host list' – typically /etc/hosts.equiv or .rhosts in user's home directory. The trusting host setup is commonly used by a suite of 'r' commands on UNIX systems. The 'r' commands include 'rsh', 'rlogin' and 'rexec'. Rlogin allows remote login. Rexec and rsh can be used to issue a command from a remote system. The trusting host scheme is convenient for the users. But if one trusted host is compromised, it opens the door for unauthorized access to other trusted hosts, without needing a password.

The other problem with telnet and rlogin is that they transmit the terminal session on the network in clear text. Attackers could snoop the network and capture everything you type on the terminal, including your login name and password.

SSH is a tool that also provides terminal access to a machine. The biggest advantage of SSH is that it encrypts your password and terminal session. You should use SSH in place of rlogin or telnet as much as possible. This is especially important for system or network administrators who often need to log onto privileged accounts.

FTP (File Transfer Protocol) is a program that allows users to transfer files from one machine to another. FTP servers use two connections – one connection at port TCP 21 for commands or control, and another connection at port TCP 20 for data. Normally, an FTP session is established as the following. An FTP client connects to a server, running at port 21. The FTP client informs server which port it is using, via the "PORT" command. When the client issues a command, that requires data, the server opens a TCP session on port 20 from itself to the client port, for transmitting the actual data. In this case, firewall needs to allow inbound connections on random high number ports, for transmitting data packets from FTP servers to clients on the internal network. Another implementation of FTP is called "passive FTP". With passive FTP, a client informs the server of the port number it's using, via the "PASV" command. In this scheme, client initiates the connections for data transfer.

Vulnerabilities were found in both "PASV" and "PORT" implementations that could allow FTP servers behind the firewall to be compromised. If you need to allow inbound FTP connections from the Internet, make sure you patch the FTP servers to the later version, and restrict access to the source IP or require strong authentication at user level.

For distributing files on the Internet, anonymous FTP is the most commonly used program. By convention, users can login to an anonymous FTP server, using the login name 'anonymous' and their email address as password. When configuring an anonymous FTP server, you should pay special attention to the ownership and

permissions of the files and directories in the FTP login area. No files or directories in the anonymous area should be writable or owned by the anonymous FTP login. Vulnerabilities were found in a number of popular Anonymous FTP tools. One example is the remote buffer overflow problem (CERT Advisory: FTP-Buffer-Overflows) that can lead to a potential root compromise. Anonymous FTP services should not be allowed into your internal network. Normally anonymous FTP server is connected to your DMZ (screened network) for access from the Internet.

Microsoft file and print shares are done through NetBIOS port TCP 139. When Windows share is not configured properly, it often leads to file corruption or information theft. The best defense is to block inbound connections to NetBIOS services on your router or firewall.

Syntax of the filter

For CheckPoint FW-1, the rules for filtering telnet, rlogin and FTP are listed in below. The NBT traffic is rejected on the firewall. Note that the data field 'Track' for NBT traffic is blank, so that the broadcast traffic will not be logged.

Firewall Policy

	SOURCE	DESTINATION	SERVICES	ACTION	Track
	Any	Any	NBT	Reject	
	Mynet	Not Mynet	FTP rlogin telnet ssh	Accept	Long

Network Objects

Name	Type	Location	IP Address	Netmask
Mynet	Network	Internal	192.168.1.0	255.255.255.0

Service Objects

Name	Type	Port	S_port from	S_port To	Match	Prolog	members
NBT	Group	-	-	-	-	-	Nbsession Nbname nbdatagram

Description of each of the parts of the filter

An FW-1 rule consists of the following data fields - Source, Destination, Service, Action and Track. The data field Action can have one of the following values - accept, deny, reject, authenticate or other user-defined actions. Track allows you to specify which type of log tracking is required.

The first rule in the firewall policy indicates that all NetBIOS traffic should be dropped. The second rule indicates that FTP, rlogin, telnet and ssh are allowed from the internal network to the Internet.

Mynet is defined as a Network Object in the FW-1. The IP range of Mynet is 192.168.1.0 and the netmask is 255.255.255.0.

NBT is a Service Group which is pre-defined in the "Service Objects" in FW-1. The NBT group includes the protools nbdatagram (port 138), nbname (port 137) and nbsession (port 139).

Explain how to apply the filter

To define a network object, you can use Network Object Manager from the FW-1 GUI. The command to start the GUI is fwui. Choose Network Objects from the Manage menu. Click New and a menu will appear that allows you to select the type of objects. For the object "Mynet" in this example, the object type "network" was selected.

To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu. To modify a rule, right click on the value of the data field you want to change. A menu will be displayed from which you can choose the new value.

Most firewalls allow more outbound services than the login services listed in this example. The purpose of this example is to demonstrate the filter rules for the specified login services on an FW-1.

Explain how to test the filter

Once the firewall policy is installed, you can test the outbound services, by telnet'ing to an address on the Internet, for example. The outbound terminal services (telnet, rlogin, etc) are logged on the firewall. The connection logs can be viewed via FW-1 log viewer.

Assignment 3 – RPC and NFS – Portmap/rpcbind (111/tcp and 111/udp), NFS (2049 and 2049/udp), lockd (4045/tcp and 4045/udp).

Introduction

RPC – Sun Microsystems Remote Procedure Call protocol supports many network services. NIS (Network Information Service), NFS (Network File Systems), and some implementation of X11 are examples of services based on RPC.

Multiple vulnerabilities exist in RPC. Using the command 'rpcinfo', for example, an intruder could obtain information including the RPC services that are being run, the port numbers and the locations of the programs. The information could be extremely useful when developing attacks. RPC is designed with local Ethernet connectivity in mind. It is dangerous to open RPC on the firewall.

NFS mount's access control is generally done using UNIX login ID or group ID. This security mechanism works fine for accessing files on a local network. It is not suitable to

make NFS mounts available over the Internet. There is also a risk that someone could have setuid'ed programs in the NFS mounted file systems that can pose risks to client machines. NFS servers, running certain implementations of 'mounted', are vulnerable that remote users could gain administrative access to the NFS file server. 'Mounted' is a RPC service that handles requests to use (mount) a file system from a client. NFS services should not be installed on the firewall. NFS services should also be blocked by the firewall.

Syntax of the filter

For CheckPoint FW-1, the RPC and NFS services, or any other unauthorized services can be blocked by the default 'deny all' rule. The 'deny all' rule is the last rule in the rule base.

Firewall Policy

	SOURCE	DESTINATION	SERVICES	ACTION	Track
	Any	Any	Any	Drop	Long

Description of each of the parts of the filter

RPC and NFS services are blocked by the firewall. By default firewall drops all unauthorized traffic but does not log the denied traffic. The purpose of this rule is to drop and log unauthorized traffic, in this case, RPC and NFS requests.

RPC is required for CDE and Openwin to run on a Sun Solaris system. If you need to use a windowing system for remote access, instead of allowing RPC on the firewall, a good alternative is use SSH.

Explain how to apply the filter

First start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

Explain how to test the filter

The testing of the NFS and RPC filters can be done by probing the ports on the firewall using one of the freely available tools on the Internet. Unauthorized accesses will be logged and can be found in the log files.

Assignment 4 – NetBIOS in Windows NT – 135 (tcp and udp), 137 (udp), 138 (udp), 139 (tcp). Windows 2000 – earlier ports plus 445 (tcp and udp).

Introduction

Windows NT functions use ports 135 through 139 and Windows 2000 use ports 445 plus ports 135 to 139. The NT functions include NT User Manager, Server Manager, Event Viewer, Registry Editor, NT diagnostics and Directory replication. The NT functions are based on a trust based environment that presents a significant risk. There were numerous bug reports and vulnerabilities found in NetBIOS. NetBIOS services are a source of great security danger. In many cases, intruders could gain full control of your NT server from remote, crack password or crash the services on server. The best defense is to block inbound connections to NetBIOS services on your router or firewall.

Syntax of the filter

Firewall Policy

	SOURCE	DESTINATION	SERVICES	ACTION	Track
	Any	Any	NBT	Reject	
	Any	Any	Any	Drop	Long

Description of each of the parts of the filter

The first rule indicates to the firewall that all NBT packets should be rejected, and no log will be kept for the denied packets. The rest of NetBIOS protocols will be blocked by the 'deny all' rule at the end of the rule base.

Explain how to apply the filter

To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu. To modify a rule, right click on the value of the data field you want to change. A menu will be displayed from which you can choose the new value.

Explain how to test the filter

You can verify that the NetBIOS services are blocked by the firewall by scanning the ports using one of the freely available tools on the Internet.

Assignment 5 – X Windows – 6000/tcp through 6255/tcp.

Introduction

X windows allows an application to run on one system and display on a different machine. X windows is a wonderful system and provides great convenience for the users. But from security's point of view, X server gives away control of the user's terminal, mouse and keyboard. For example, an application that has connected to the X11 server could obtain full control of the keyboard and intercept your password. An

intruder could also connect to your X server and then make screen dumps or detect your keystrokes.

A number of security mechanisms are available to protect X server. The first protection is host-based authentication using a command called 'xhost'. It allows a machine to connect to the X server only if it is listed in its host list. But xhost does not provide authentication at the user level. In other words, once a machine is given permission to connect to an X server via xhost, any users from that machine can connect to that X server. The second method authenticates using a token called 'magic cookie". In this case, only those applications that share the same 'cookie' can connect to the server. There are also other security methods available, such as secure RPC. However, RPC also has some vulnerabilities of its own (see "Assignment 3" on RPC). In general, even with the protection through the use of authentication methods, X windows is still dangerous.

Syntax of the filter

On CISCO IOS, the filters to block X windows is the following.

```
access-list 101 deny tcp any any eq 6000 log
access-list 101 deny tcp any any eq 6001 log
```

Description of each of the parts of the filter

The filter denies all packets with the port number 6000 or 6001. The filter also tells router to log the denied packets.

In the example of the X11 filter, access-list 101 is a Cisco "Extended" access-list. The list number specified in an Extended access list must have a value between 100 and 199. Unlike Standard access list that it filters packets based only on the source address, Extended access lists allow more parameters. In addition to source address, Extended access list allows to filter traffic based on destination addresses, and specify protocol types or protocol numbers.

X11 uses port 6000 for the first X server on the system. The second server is at 6001 and so on. A common approach is to block a range of these ports. In the example of X Windows filters, ports 6000 and 6001 are blocked.

Explain how to apply the filter

The filter is applied to the Cisco router by entering the commands as shown in the example above at the router's command line interface.

Explain how to test the filter

The testing of the filter can be done by testing X windows connections from a host on the Internet. Since the denied accesses will be logged. The denied packets can be found in the log file.

Assignment 6 – 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).

Introduction

DNS (Domain Name System) is probably the world's largest distributed database in use. DNS is a client/server system used to map host names to IP addresses and vice versa. Each DNS domain is administered by a name server. Domain updates are done on the primary server. New DNS maps are propagated from primary server to the secondary servers via zone transfer. On the client side, the name service is called "resolver". "Resolver" queries DNS server to resolve names. A DNS client makes connections to DNS servers on port UDP 53. Port TCP 53 is usually used for zone transfers.

Different tools, such as nslookup, can be used on the client for retrieving information from DNS. At the same time, nslookup can also be used for reconnaissance purposes. An intruder could obtain the host information, IP addresses, and learn about your network simply through DNS lookups. The key steps to secure DNS are to limit DNS queries and to restrict zone transfers to authorized IP addresses only, and log zone transfer attempts.

Most firewalls are designed to use split DNS. By split DNS, I mean an internal DNS for resolving internal names and an external DNS for resolving external names. The internal host names are hidden behind the firewall. Only users on the internal network can query the internal DNS. The external DNS maintains those entries needed for external services, such as the hostname and IP address of your external Web server. The external DNS can not query the internal DNS. The internal DNS can look up the external DNS. Finally, Zone transfers are allowed only between external DNS and the secondary domain servers.

A number of vulnerabilities have been found in the earlier releases of the Berkeley Internet Name Domain (BIND). BIND is the most widely used implementation of DNS. It is critical that you apply patch or update to a later version of BIND. Intruders could exploit the vulnerabilities to gain root access on your system, disrupt the normal operation of your name server, or cause the name server to crash. Another example is DNS cache poisoning. In the case of a contaminated DNS cache, the address of a targeted web server could be falsely re-directed to the attacker's web site.

LDAP (Lightweight Directory Access Protocol) is a lightweight version of the X.500 directory access protocol. A LDAP client makes a TCP connection to a LDAP server listening on port 389, over which it sends requests and receives responses. LDAP provides access to a central database containing user level security information that can be shared among multiple network applications. You should guard your LDAP database carefully.

Syntax of the filter

For FW-1, the filter rule is the following.

In this example, the primary DNS server, named "DNSserver", resides in the DMZ (or the screened network). The secondary DNS server (named SecondDNS) is one of our ISP's systems, and is on the Internet.

Firewall Policy

Rule	SOURCE	DESTINATION	SERVICES	ACTION	Track
1	DNSserver	Not Mynet	Domain-udp	Accept	Long
2	Any	DNSserver	Domain-udp	Accept	Long
3	DNSserver	Any	Domain-tcp	Accept	Long
4	Any	DNSserver	Domain-tcp	Accept	Long

Network Objects

Name	Type	Location	IP Address	Netmask
DNSserver	Host	External	199.184.12.6	255.255.255.0
SecondDNS	Host	External	128.204.15.7	255.255.0.0
Mynet	Network	Internal	192.168.1.0	255.255.255.0

Description of each of the parts of the filter

Rule 1 - Allows DNS queries and answers (port UDP 53) from our primary DNS server (199.184.12.6) to DNS servers on the Internet.

Rule 2 – Allows DNS queries and answers (port UDP 53) from both the internal network and the Internet to the primary DNS server.

Rule 3 – Allows DNS queries and zone transfers (port TCP 53) from the primary DNS server to the Internet.

Rule 4 – Allows DNS queries and zone transfers (port TCP 53) from the Internet to the Primary DNS.

We have 3 items defined in the Network Objects table. They are DNSserver, SecondDNS and Mynet . The object type of "DNSserver" is "Host". Its IP is 199.184.12.6 and netmask 255.255.255.0. DNSserver is installed in the screened network, and therefore the "Location" is External. SecondDNS's host IP is 128.204.15.7 with the netmask 255.255.0.0. Mynet is the internal network, which is a class-C network. Mynet is defined as a "Network" object in FW-1. The IP range of Mynet is 192.168.1.0 and the netmask is 255.255.255.0.

Rule 3 and 4 allow zone transfers and DNS queries between the DNSserver and the Internet. It is important to ensure that your primary DNS server is limiting zone transfers to only those Secondary name servers. In the BIND 8.x configuration file - /etc/named.conf, zone transfer can be limited using the command allow-transfer.

LDAP (port 389) is not specifically allowed and therefor is blocked by the firewall.

Explain how to apply the filter

First start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

Explain how to test the filter

The testing of the DNS rules can be done using the command nslookup. You can query DNS via the command 'nslookup'. Zone transfers can also be initiated using the command nslookup. At nslookup prompt, type "ls –d" that requests a zone transfer from the specified domain name. If you initiate a zone transfer from a host that is not a registered secondary server, the firewall will deny the request. The denied packets will be logged and can be found in the log files.

Assignment 7 – Mail – SMTP (25/tcp) to all machines, which are not external mail relays, POP (109/tcp and 110/tcp), IMAP (143/tcp).

Introduction

Internet mails are usually transported using a protocol called SMTP (Simple Mail Transport Protocol). SMTP allows the sender to specify a return address in the "MAIL FROM" command. Based on SMTP, you do not know if mail actually came from the sender, as the mail header suggests. There are a number of vulnerabilities associated with SMTP. SMTP could be used for DOS (Denial of Service) attack. Someone could flood a mail server with tons of messages and with bogus FROM addresses. There is no good way for you to find out where the mail came from. It is worth noting that sendmail, a widely used implementation of SMTP installed in most UNIX systems, has a long history of security vulnerabilities. A serious vulnerability found in earlier versions of sendmail allows a remote user to execute arbitrary commands on the local system with root privilege. In most cases, you should avoid running sendmail on the firewall.

Mail gateway can ensure that outgoing mails have correct FROM addresses. While this does not prevent your network from receiving emails with bogus headers, it can prevent propagating mails that do not conform to the standard. Mail gateway should also deny incoming mail with mail relay. Mail relay occurs when a mail server processes an email where neither the sender nor recipient is a local user. By permitting mail relay, you risk your mail server to DOS attacks, and at the same time, you are letting spammers use your systems for relaying junk mail to the Internet.

POP and IMAP are mail protocols that allow users to access their emails from remote systems. Serious vulnerabilities that allow remote users to execute arbitrary commands with privileged accounts, were found in some implementations of IMAP or POP servers.

Syntax of the filter

For FW-1, the filter rule for SMTP is the following.
In this example, Mailserver is in the DMZ network.

Firewall Policy

Rule	SOURCE	DESTINATION	SERVICES	ACTION	Track
1	Mailserver	Any	SMTP	Accept	Long
2	Any	Mailserver	SMTP->smtpfilter	Accept	Long
3	Mynet	Mailserver	POP IMAP SMTP	Accept	Long

Network Objects

Name	Type	Location	IP Address	Netmask
Mailserver	Host	External	199.184.12.5	255.255.255.0

Description of each of the parts of the filter

Rule 1 – Allows outgoing mail from the SMTP Mail server (Mailserver) to the outside world.

Rule 2 – Allows incoming mail from the outside world to the Mailserver. The incoming messages are filtered by FW-1 SMTP security server, smtpfilter in this example. Smtpfilter can be defined in the "Security Servers" of the "Control Properties" from FW-1 GUI. One of the features provided by FW-1 SMTP Security Server is that it filters messages with the header in the format of user%host.com@relayhost.com. To ensure that your mail server does not allow mail relay, you should also install a later version of the mail server software (Sendmail, for example) on the mail host.

Rule 3 – Allows POP and IMAP access from mail clients on the internal network (Mynet) to the Mailserver.

In the Network Objects table, "Mailserver" is "Host". Its IP is 199.184.12.5 and netmask 255.255.255.0. Mailserver is installed in the DMZ (the screened network), and therefore the "Location" is External.

Explain how to apply the filter

First start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

Explain how to test the filter

The testing of the filter rules can be done by sending emails to the mail server host from the Internet. To test if mail relay is disabled, you can send an email with the TO address in the format of user%host.com@Mailserver.com, where mailserver.com is the address of your mail server. IMAP and POP are not allowed from the outside to the mail server. Failed attempts will be logged and can be found in the log files.

Assignment 8 – Web – HHTP (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.).

Introduction

HTTP (Hypertext Transfer Protocol), the World Wide Web protocol is a TCP based service. Most HTTP servers use port TCP 80, but sometimes they run at high number ports (greater or equal to 1024). The ability to provide HTTP services on high number ports allows users to start up their own web servers. For example, a user could run a web server on a host named "webserver.com" and randomly choose a port number, such as 8080. The Web service could then be accessed via the URL, http://webserver.com:8080. This is convenient for the users, but it poses security dangers and complicates the configuration of a packet filtering firewall, if you have to allow outbound HTTP access through random high number ports.

Many Web servers use CGI programs to allow services such as Web forms for data collections or on-line shopping cart applications. A number of vulnerabilities were found on Web CGI applications. Vulnerable CGI programs could lead to unauthorized access to the systems, credit card number theft or vandalize web pages. It is important that you carefully configure the security of a Web server to restrict the server's operations to a certain hierarchy of your server's file system.

HTTP clients (Web browsers) are susceptible to a number of security vulnerabilities. When clients visit a Web server, the data returned by the server can come in all sorts of formats. They can be video, audio, or HTML with imbedded scripting commands, for example. With the data received from Web servers, browsers will perform the necessary operations; play the audio, display the pages, run the scripting commands or invoke a program outside of the browser on local system. For example, a browser will load Adobe Acrobat Reader if the browser sees a PDF format file. This scheme presents security problems. Intruders could take advantage of the external program invoked by the browser, or cause a vulnerable program to be loaded on your system from browser.

Another item worth noting here is that HTTP proxy is a very popular method for handling browser requests from internal network. HTTP clients connect to the proxy server using a fixed port. The request is then sent from the proxy to the destination Web server using either a standard port or arbitrary ports. Proxy servers can provide a feature

called page caching that dramatically improves browser performance and reduces network traffic. Some web proxy services also provide content filtering that blocks unwanted Web sites.

Secure HTTP services run at port TCP 443. Secure HTTP provides authentication and encryption. It is designed to ensure privacy for the information passed over the Internet.

Syntax of the filter

For FW-1, the filter rule for HTTP and HTTPS is the following. The following assumes that an HTTP proxy server (named Webproxy) is configured to provide the HTTP proxy services for users on the internal network. Clients on the internal network connect to the HTTP or HTTPS servers on the outside through the HTTP proxy.

Firewall Policy

Rule	SOURCE	DESTINATION	SERVICES	ACTION	Track
1	Any	Webserver	HTTP HTTPS	Accept	Long
2	Webproxy	Any	Any	Accept	Long

Network Objects

Name	Type	Location	IP Address	Netmask
Webserver	Host	External	199.184.12.7	255.255.255.0
Webproxy	Host	Internal	192.168.1.10	255.255.255.0

Description of each of the parts of the filter

Webserver is in DMZ. Webproxy – HTTP proxy server is on the internal network.

Rule 1 – Allows HTTP or HTTPS connections from both the internal network and the Internet, to the external Web server.

Rule 2 – Allows the HTTP Proxy server (Webproxy) to connect to the HTTP servers on any machine on the Internet. By specifying "Any" in the Service data field, the rule allows outgoing connections from Webproxy to the HTTP servers on the Internet on non-standard ports (not port 80). Note that for FW-1, to enable users to use FTP from their browsers, you need to open FTP from the internal network to the outside.

Explain how to apply the filter

First start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

Explain how to test the filter

The testing of the HTTP/HTTPS filter rules can be done by using a web browser. You can connect from the internal network to the HTTP or HTTPS servers on the outside. You should also test incoming HTTP or HTTPS request from the Internet to your external web server, Webserver in this example. The testing of the filters should also include connecting to HTTP servers that use non-standard ports.

Assignment 9 – "Small Services" – ports below 20/tcp and 20/udp, time (37/tcp and 37/udp).

Introduction

Small Services are ports TCP and UDP below 20 and 37. These services are rarely used. Only those services that are specifically authorized should be allowed on the firewall; these unused ports should be blocked by the firewall.

Although the services are seldom used, various attacks have been targeted at the Small Services. Port 0 is invalid. Both port 0 and port 1 (tcpmux) could be scanned to identify the Operating Systems in a network mapping. Port 7 (Echo) and 19 (Chargen) could be used for DOS attacks; "echo-loop" or the "fraggle" DOS attack are examples of such attacks.

Syntax of the filter

For CheckPoint FW-1, the Small services, or any other unauthorized services can be blocked by the default 'deny all' rule. The 'deny all' rule is the last rule in the rule base.

Firewall Policy

	SOURCE	DESTINATION	SERVICES	ACTION	Track
	Any	Any	Any	Drop	Long

Description of each of the parts of the filter

The filter rule blocks any unauthorized services. The denied packets are logged and can be found in the log files.

Explain how to apply the filter

First start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

Explain how to test the filter

The testing of the filter can be done by scanning the ports from the outside using one of the tools available from the Internet. Since the services are blocked by the firewall, the denied packets will be logged and can be found in the log files.

Assignment 10 – 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).

Introduction

TFTP (Trivial File Transport Protocol/port UDP 69) supports simple file transfer. The protocol is often used by systems to download boot code. Since there is no authentication in TFTP, if the server is misconfigured, intruders could gain access to any files on the system.

Finger is a utility that allows remote users to query user information from a system. The protocol has a number of security issues. Intruders could find out what Operating Systems are being used or obtain user information. There are also known "buffer-overflow" bugs associated with finger.

NNTP (Network News Transfer Protocol/port TCP 119) is used for posting, reading, and exchanging USENET news between NNTP servers. Access to the NNTP server should be restricted to certain IP addresses. An unrestricted NNTP server allows post and read from anyone.

NTP (Network Time Protocol/port TCP 123) is used to synchronize a system's clock with the NTP servers on the Network. Keeping the clocks on your systems in sync is extremely important. It allows you to match the logs from different machines. The NTP sever itself could be a target for various attacks. Attackers could alter the time on a system. Some authentication systems that are time based could be exploited to gain unauthorized access.

LPD (Printing Protocol/port TCP 515) provides printing services. Modern printers have built-in ethernet cards with its own disks. In theory, it's possible to download a printer command file to a printer and cause disruptions to the printing service. This could be a form of DOS attack.

Syslog is a UNIX utility that allows logs to be collected on a central log server. Keeping logs in one place gives you a global view of what might be happening on your network. Many network devices now use syslog to report problems and status. Cisco router is one example. Constant monitoring of the logs is important.

SNMP (the Simple Network Management Protocol) is widely used by network managers to administer and monitor all types of network devices. SNMP manageable devices range from router, hub, and printer to computers. SNMP access is authenticated using a pair of

unencrypted 'community strings'. Most vendors ship devices to customers with the 'standard' community strings – 'public' for read access, and 'private' for read-write access. It is not uncommon to find that a network device is left running with the default community strings intact. With the known SNMP 'community strings', intruders could easily obtain all sorts of information about your network.

BGP (Border Gateway Protocol/port TCP 179) is a routing protocol that supports the Internet backbone. BGP is installed on the edge of your network. It is critical to the operation of the network. Various attacks have been developed that are targeted at BGP protocol. The areas of attacks include, SYC flooding, RST attacks, DATA injecting attacks, and "Session Hijacking" attacks. "Session Hijacking" attacks attempt to gain control of the TCP connection from the neighboring BGP server to your BGP server.

SOCKS is a protocol that tunnels traffic through firewalls. An application gateway implemented using SOCKS allows hosts behind the firewall to access the Internet through a single IP address. If the server is misconfigured, intruders could gain access through the tunnel into your internal network.

Syntax of the filter

Given an FW-1, the filter rules for the services – NNTP, syslog and NTP are the following.

Firewall Policy

Rule	SOURCE	DESTINATION	SERVICES	ACTION	Track
1	MyNNTP	OtherNNTP	NNTP	Accept	Long
2	OtherNNTP	MyNNTP	NNTP	Accept	Long
3	Mynet	MyNNTP	NNTP	Accept	Long
4	NTPhost	TimeServers	NTP	Accept	Long
5	Mgthost	ExtRouter	SNMP	Accept	Long
6	ExtRouter	Mgthost	Syslog SNMPtrap	Accept	Long

Network Objects

Name	Type	Location	IP Address	Netmask
MyNNTP	Host	External	199.184.12.8	255.255.255.0
OtherNNTP	Host	External	128.204.17.5	255.255.0.0
Mgthost	Host	External	199.184.12.10	255.255.255.0
NTPhost	Host	Internal	192.168.1.11	255.255.255.0
Mynet	Network	Internal	192.168.1.0	255.255.255.0
TimeServers	Group	-	-	-
ExtRouter	Router	External	199.184.11.1	255.255.255.0

SNMP filter ----
Given a Cisco router that routes for a class-C address space 199.184.12.*, and assuming the management station is 199.184.12.10, the SNMP filter is the following:

```
!                    --------- SNMP filters -----------
! Define the host permitted to make SNMP read requests to this router.
access-list 21 permit 199.184.12.10
!
! Give RW (Read-Write) access to SNMP to the management station.
snmp-server community secret RW 21
!
! Tell router to send SNMP messages to the specified host using
! secret community.
snmp-server host 199.184.12.10 secret
!
! Enable traps for authentication failures caused by SNMP access violation.
snmp-server trap-authentication
!
! Traps for router configuration changes.
snmp-server enable traps config
```

Description of each of the parts of the filter

FW-1 Policy -------

Rule 1 and Rule 2 – Allow USENet news both ways between my NNTP server, and my news feed server on the Internet.

Rule 3 – Allows clients from the internal network to read and post news to my NNTP server.

Rule 4 – Allows a local NTP host to talk to a group of NTP servers on the Internet to maintain an accurate source time on the local NTP host. The NTP clients on the internal network communicate with the internal NTP host to keep their clocks in sync.

Rule 5 – Allows the management hosts (Mgthost) to contact the external SNMP device (ExtRouter) for information.

Rule 6 – Allows the external SNMP device (ExtRouter) to contact the internal SNMP trap server (Mgthost) via SNMP trap service. It also allows the external SNMP device (ExtRouter) to send problems or status, to the management station via syslog.

"TimeServers" is a Group object, defined in the Network Objects in FW-1. The Group consists of multiple external NTP hosts. Each group member in the Group definition is a Host Network Object.

SNMP should be blocked for traffic coming from the outside to the internal network. In the example of the SNMP FW-1 filter rule, it allows SNMP traps and syslog messages sent from the external router to the management station. It also allows the management station to contact the external router via SNMP. On the router itself, Cisco filters are implemented to limit SNMP access to the management station only.

Finger can be allowed from the internal network to the outside. But finger should not be allowed from the outside to the internal network. In this example, finger is blocked by the firewall.

BGP is the routing protocol for the border router. The packets should not be allowed through the firewall from outside to the internal network.

SOCKS is blocked by the firewall.

LPD is blocked. You will not be able to print from a system outside of the firewall to a printer on the internal network.

TFTP is blocked by the firewall. If you have a device outside of the firewall, that needs TFTP for downloading boot code, you could either set up a bastion host for booting or choose other boot methods, such as diskettes.

The NNTP server is in the screened network (DMZ). The filter rule allows the NNTP server to exchange USENet news with the neighboring NNTP news servers on the Internet. Clients on the internal network are allowed to access the NNTP server to read or post news.

Explain how to apply the filter

For FW-1 filter rules, first start up FW-1 GUI. To add a rule from FW-1 GUI, select "Edit" and then choose "Add Rule" from the menu.

For Cisco, the filters are applied to the Cisco router by entering the commands as shown in the examples above, at the router's command line interface.

Explain how to test the filter

The testing of SNMP filters can be done by sending SNMP messages to the router. Denied packets can be found in the logs. The testing of the filter rules can be done, using the tools freely available from the Internet, by scanning the ports on the firewall. To test if those services allowed by the firewall work correctly, you can use the appropriate client applications, such as NTP client, or NNTP newsreader.

Assignment 11 – ICMP – block incoming echo request (ping and Windows traceroute), block outgoing echo replies, time exceeded and unreachable messages.

Introduction

ICMP is an IP protocol designed to deliver simple control messages, requests and replies from one machine to another. For example, ICMP can be used to detect if a host is alive, or to inform another host of an unavailable service. However, quite often the protocol has been manipulated for malicious attacks.

Compared to other IP protocols, ICMP is unique in a number of ways. ICMP does not use port number. It uses message type and code, to identify a particular ICMP function. ICMP does not have a notion of client/server or reliable packet delivery. It does not require a response in some cases. Also, ICMP supports broadcast. An ICMP message can be sent to multiple hosts at one time.

ICMP echo request is a most common technique for network mapping. An ICMP echo request message can be sent to one host, or sent via broadcast to a number of hosts. By using ICMP messages, attackers could find out the IP addresses of live hosts in a network. Once the IP addresses are found, attacks can be targeted at those hosts. Since most firewalls block incoming ICMP echo requests, techniques were developed to do network mapping using other ICMP messages types, such as mapping using ICMP netmask requests. ICMP echo messages, request or reply, can be used for malicious activities, and not just network mapping. Attackers use ICMP to launch DOS attacks, and to tunnel commands in an ICMP packet. Smurf , WinFreeze attack and TNF attack are examples of DOS attacks. Loki is an example that attackers use ICMP as tunneling mechanism for malicious activities.

Syntax of the filter

Given a Cisco router that routes for a class-C address space 128.204.*.*, the following is the filter that allows only specific ICMP packets from the outside.

```
! Allow only specific ICMP messages.
!
!  net unreachable
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 3 0
! host unreachable
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 3 1
! port unreachable
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 3 3
! packet too big
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 3 4
! administratively prohibited
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 3 13
! source quench
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 4
! TTL exceeded
access-list 101 permit icmp any 128.204.0.0 0.0.255.255 11 0
```

Given a Cisco router that routes for a class-C private address space 192.168.1.*, the filter that blocks outbound ICMP messages is the following:

Access-list 102 deny icmp any any log

Description of each of the parts of the filter

In the example of ICMP filters, access-list 101 is a Cisco "Extended" access-list. The list number specified in an Extended access list must have a value between 100 and 199. Unlike Standard access list that it filters packets based only on the source address, Extended access lists allow more parameters. In addition to source address, Extended access list allows to filter traffic based on destination addresses, and protocol types or protocol numbers. In the example of ICMP filters, only certain ICMP messages are allowed by the router. The allowed ICMP messages include, destination unreachable, packet too big, administratively prohibited, source quench, and TTL exceeded. The intent is to permit only those ICMP messages needed by the router. The most commonly used ICMP echo request and reply messages are blocked by the router.

The ICMP messages, net unreachable, host unreachable, and port unreachable, inform the sending hosts that it can not route or deliver packets to the destination. A "Packet too big" ICMP message occurs when a datagram is too big and fragmentation is needed, but the Don't Fragment (DF) bit is set. The administratively prohibited message is sent when a sending host is denied access to a service. In this case, it could be that the port or the protocol is blocked or the IP is not allowed access. ICMP source quench message is sent when packets are arriving at the router too quickly for it to process. TTL (Time To Live) exceeded message is used to tell the sending host that the 'hop count' has reached zero and the packet will be discarded. This is to prevent a routing loop.

The ICMP filter for the outbound traffic, indicates to the router that from any source address to any destination address, all ICMP messages are blocked and logged. If you need to allow ping from certain internal hosts to the Internet, the filter rule can be changed to only allow ping for a group of internal IP addresses.

Explain how to apply the filter

The filters are applied to the Cisco router by entering the commands as shown in the examples above, at the router's command line interface.

Explain how to test the filter

The testing of the ICMP filters for packets from the outside can be done using a number of commands or tools freely available on the Internet. It's simply a matter of generating a packet with the ICMP message types, and send it from outside. Tests should also be done to verify those denied ICMP message types are indeed blocked by the router. Since ICMP echo request is not permitted, the easiest test is to 'ping' a device behind the

router. Other ICMP messages that are denied by the router, such as re-direct and address mask request messages, should all be tested.

The testing of the ICMP filter for packets sent from the internal network, can be done by using the command 'ping' or tools available from the internet. Internal hosts are not allowed to send ICMP message, the denied packets can be found in the logs.

Appendix: Firewall Rules – Putting it together.

In this document, we have shown the ingress, egress, ICMP, SNMP and X11 filters on Cisco IOS. Using the Cisco filter rule is a simple and efficient first line of defense before the inbound traffic can reach your firewall. On the firewall, the packets are examined by the firewall filter rules. The FW-1 firewall filter rules presented in this document are listed as the following. Please note that in addition to the rules given in the examples, two new filter rules were added – rule 1 and rule 3. Rule 1 is needed to allow firewall administrator connect to the firewall itself. Rule 3 is a "Firewall Lock Down" rule. The FW-1 firewall filter rules are numbered in the order they are applied on the firewall.

Summary of the firewall filter rules.

Rule	SOURCE	DESTINATION	SERVICES	ACTION	Track
1	FW-Admin	Firewall	Firewall-1	Accept	Long
2	Any	Any	NBT	Reject	
3	Any	Firewall	Any	Drop	Long
4	Any	Webserver	HTTP HTTPS	Accept	Long
5	Mynet	Mailserver	POP, IMAP,SMTP	Accept	Long
6	Mailserver	Any	SMTP	Accept	Long
7	Any	Mailserver	SMTP->smtpfilter	Accept	Long
8	Webproxy	Any	Any	Accept	Long
9	Mynet	Not Mynet	FTP rlogin telnet ssh	Accept	Long
10	Mynet	MyNNTP	NNTP	Accept	Long
11	DNSserver	Not Mynet	Domain-udp	Accept	Long
12	Any	DNSserver	Domain-udp	Accept	Long
13	DNSserver	Any	Domain-tcp	Accept	Long
14	Any	DNSserver	Domain-tcp	Accept	Long
15	NTPhost	TimeServers	NTP	Accept	Long
16	Mgthost	ExtRouter	SNMP	Accept	Long
17	ExtRouter	Mgthost	Syslog SNMPtrap	Accept	Long
18	MyNNTP	OtherNNTP	NNTP	Accept	Long
19	OtherNNTP	MyNNTP	NNTP	Accept	Long
20	Any	Any	Any	Drop	Long

Rule 1 - This rule has to go before the "Firewall Lock Down rule" which is rule 3.

Rule 2 - NetBIOS is very noisy and NBT broadcast generates a lot of traffic. This rule is used often enough to be placed before the Firewall lock down rule.

Rule 3 - Firewall lock down rule needs to be placed before the Internet access rules.

Rule 4 – This rule allows access to the external web server from the Internet and the internal network. Rule 4 to 10 are essential services, provided to the internal network and the Internet. They are arranged in the order based on how frequent the rules are used.

Rule 5, 6, and 7- Email is one of the most important network services. The rule will be used quite often. The Email services also generate a lot of traffic.

Rule 8 – This is the HTTP proxy server for the internal network. Similar to Email, this service is frequently used and generates a lot of traffic.

Rule 9 and 10 - These are terminal access and USENET news access from users on the internal network.

Rules 11, 12, 13 and 14 – These are DNS services. Every network access potentially needs to query the DNS. Except the zone transfers filter rule, the DNS rules are used fairly frequently.

Rule 15, 16 and 17 – These are important network management services. Although they don't normally generate a lot of traffic, the rules could be used more often than rule 18 and 19.

Rule 18, 19 – News feed transfers are usually run at night in a batch job.

References

Help Defeat Denial of Service Attacks: Step-by-Step
http://www.sans.org/dosstep/index.htm

Lance Spritzner Whitepapers & Publications
http://www.enteract.com/~lspitz/papers.html

Firewall-1 FAQs
http://www.phoneboy.com/fw1/

Understanding the FW-1 State Table
http://www.enteract.com/~lspitz/fwtable.html#table

Auditing Your Firewall Setup
http://www.enteract.com/~lspitz/audit.html

Configure and use Perimeter routers
http://pasadena.net/cisco/secure.html

Security Portal Weekly Check Point Security Digest
http://www.securityportal.com

Internetworking with TCP/IP, *D. Comer, Prentice Hall*

Firewalls and Internet Security, *W. Cheswick & S. Bellovin, Addison-Wesley*

Building Internet Firewalls, *D. Capman & E. Zwicky, O'Reilly & Associates, Inc.*