



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 VPNs

GCFW Practical Assignment

Version 1.5e

Ralph McAvoy

Question #1: Security Architecture:

Introduction:

GIAC Enterprises and Joes Fortunes, a recently acquired company, are connected to the Internet. The connection between each of the companies and their ISP is a T1. Each company also has been allocated a class C range of IPs for their use. Since they have been hearing many news stories about the dangers of the Internet the management has announced their full support for an Information Technology security strategy for protecting the information resources of the company and encouraged the workers to support and implement it as well.

Business Overview:

To design an effective security architecture the company's business operations, functions, and data flow must be understood. GIAC Enterprises has a number of business relationships with suppliers, customers, business partners, a recently acquired/merged operation, and of course the public in general. In addition to these external relationships there are many internal groups and functions that must be considered when developing an integrated security architecture. Each of these groups and their requirements must be examined so as not to create an architecture that is too restrictive or one that is too loose. If the architecture is designed to be extremely restrictive the group may not be able to perform the business function or the effort to obtain the information required may become too costly. Making the security architecture too loose leaves the information technology infrastructure of the company open to compromise and potential theft of company information and proprietary data. For the purpose of this architecture six basic groups have been identified, the newly acquired company Joes Fortunes, the existing operations of GIAC Enterprises, customers, authors/suppliers, business partners, and the general public who will visit the GIAC Enterprises website out of curiosity or will visit the site to find out about GIAC Enterprises for investment purposes or business opportunities.

Group requirements:

Each of the defined groups has specific requirements relative to the function they provide for the company. Starting with the groups which are most trusted and going to the groups which are least trusted the security requirements for each group will be covered in this section.

GIAC Enterprises:

This group consists of the original company and is the most trusted of all the groups and consists of the corporate network infrastructure and the company users. These users will need to have access to the appropriate corporate network resources, such as people in finance will need access to the financial data bases while people in sales will need access to the sales data bases. All of the smaller groups within the GIAC Enterprise group have full permissions to access the Internet, http, ftp, telnet, real audio, etc. for performing their business functions, such as accessing financial institutions for money transfers or accessing websites on the Internet for research. The business use policies will be considered outside of the scope for this security architecture.

Joes Fortunes:

The merger/acquisition between GIAC Enterprises and Joes Fortunes was a friendly merger so there are no hard feelings within either of the two companies about merging into one larger company. Joes Fortunes is in a geographically different location than GIAC Enterprises with an existing network infrastructure. The details of how the two network infrastructures will be integrated internally or have been merged is also outside of the scope of this paper, but how the organizations within Joes Fortunes will gain access from their location to the functions and data at the main GIAC Enterprises campus will be covered. Since the internal groups within Joes Fortunes have been incorporated into the functions of GIAC Enterprises they will be considered to be at the same level of trust as the people in GIAC Enterprises.

Customers:

The customers are the organizations that purchase the fortunes that are created by GIAC Enterprises. This group requires access to the programs and machines that allow them to download fortune files in bulk. Since the fortunes are tightly guarded secrets and one of the main assets exchanged between GIAC Enterprises and their customers the transferred fortunes need to be protected. The process used by the customer to obtain the fortunes is to login to the GIAC Enterprises fortune server using SSL, select the fortune files they desire, and then download them. This group only requires access to the fortune distribution machine and the machine used for servicing their account. They do not need access to any of the internal machines of the GIAC Enterprises. All of their activity can be performed via the SSL/https connection.

Authors/Suppliers:

The suppliers are similar to the customers since they need to access the fortune distribution machine. The difference is that instead of downloading the fortunes the suppliers are uploading the fortunes that have been developed. Like the customers they will login and upload the fortunes they have created via SSL.

Business Partners:

The business partners are special customers. They require access to the fortune distribution machine in order to download fortunes that they in turn translate and resell. In addition the business partners require access to a machine that is used for recording

the number of fortunes they sell. Although this information will eventually be recorded in the GIAC Enterprises accounting system, the business partners do not need access to the accounting system only to a machine where they can record their sales. Their connections will also be via SSL.

General Public:

This is the group of totally untrusted visitors to the GIAC Enterprises web site. The general public will only have access to the public website and be prohibited from the access to the fortune databases or any machine used for recording sales or fortune download information. The connections to this machine will be via http only.

Architecture Rationale and Approach:

With an understanding of the types of users accessing the networks and functions or services provided by GIAC Enterprises it is now possible to begin laying out and defining the security architecture. When it comes to protecting a company's assets and resources there are many concepts and strategies that must be kept in mind, one of which is that no single defense or strategy will be sufficient to totally protect a company's networks. As a result there should be a variety of strategies used in order to provide the protection required. The strategies that will be used for this design will be based on the strategies of defense in depth, choke point, and a fail-safe stance [1].

The strategy of a defense in depth is based on the ideas that no one system by itself will provide total security but by having many layers of security it will become more costly and risky for an intruder to infiltrate the system. These layers will also provide backup and redundancy in the event of a failure at a higher layer providing additional time if a compromise does occur to detect it and react to the intrusion. The components of a defense in depth consist of network security, host security, human security, and physical security. Physical security deals with the buildings and rooms where the network and system components are kept. Access to these areas needs to be controlled and monitored so only authorized personnel can gain access to the routers, servers, switches, firewalls, and other equipment required for providing the GIAC Enterprises services.

Human Security:

Preventing security holes from being introduced into the system as a result of misconfigurations or misunderstandings about the purpose of a configuration parameter requires that all system administrators must have proper training in the configuration of the components they will be maintaining. In addition all users of the system will be instructed in the proper use of the systems and provided the requirements for password change intervals and selection. These users will also be required to run current anti-virus software using current virus definitions daily. These are the human security aspects of a layered security architecture.

Host Security:

To provide the host security layer for this architecture will require defining how each of the hosts in the system will be used and then only provide those services that are

required on each of the hosts. Once the services have been identified the OS can be loaded, all current applicable patches applied, and then the additional applications required can be loaded along with their patches and enabled. It is important that only the applicable patches be added. It is not necessary to apply the patch for a service or application that has not been loaded on the host. Adding this patch may even weaken the host security. Within the scope of this architecture we know that DNS, web, database, syslog, and mail servers will be required. In addition there will be hosts being used as firewalls. It is also important that the hosts be configured before they are connected to the network. By doing this you can ensure that a "lucky" intruder does not accidentally stumble upon an unprotected host while it is being configured. Along with being careful to only configure the required services on the host and having the most up to date patches for all software running on the host, it is important to run tripwire on the host prior to deploying it. This will provide the initial system snapshot to use in the periodic tripwire runs for determining if any unexpected changes have occurred to the baseline load. Although routers are normally not considered as hosts within a system, the configuration and the applying of patches to the routers is just as important and the same guidelines used for the host configuration need to be used for the routers within the system. Along with running tripwire it is also important to make a backup of the system and keep it in a safe place so that you have a copy of the baselined system if it ever needs to be restored. Performing all of these tasks prior to attaching the system to the network provides a fairly high level of assurance that the system is configured as planned. Copies of the router configurations also need to be captured and maintained offline for emergency recovers when necessary. The number of applications running on each of the servers will be limited providing an environment that is easier to watch and identify abnormalities when they occur. The main types of servers as specified above are DNS servers, Web servers, database servers, syslog servers, and mail servers. Sun Solaris will be the platform used for all of the firewalls, web, DNS, and syslog servers. The mail server will be an NT running exchange. All hosts that are configured for use in the DMZ or on the public subnet will be checked to ensure that NIS+, NFS, SNMP, TFTP, and SMTP services are disabled and do not start when the machine is booted. In fact most of the services provided by inetd will be disabled. Remote access to these servers will be via the open-SSH package so telnet, rlogin, and ftp can also be disabled. Each of the servers will have a cron job that runs and reports any abnormality within the system such as the web process httpd not running.

DNS Server:

To provide some protection via obscurity we will use a split horizon DNS. The DNS servers positioned on the Public subnet will provide responses to the queries from the internet and will only provide required information for the hosts which are outside of the corporate firewall, ie: those hosts which are accessible from the internet. The DNS server inside the corporate firewall will have a complete listing of all of the zones for GIAC Enterprises. Since DNS servers are heavily bombarded with intrusion attempts and bind has been identified as a potentially vulnerable piece of software the DNS servers will be dedicated to only DNS. These servers will also be located close to the system administrators for easy physical access. By locating these servers in this manner inetd does not need to run and the inetd.conf file can be eliminated

completely. The SSHd daemon will run in case remote access is required.

Web Server:

The web server will be configured with software that will provide http and https (SSL) access to the server. Any cgi scripts that are developed will use wrappers to ensure that shell meta characters and other malicious strings are stripped before execution of the script occurs. Web authors will update the pages they maintain utilizing an SSL connection and web update software GIAC Enterprises has purchased with the web server. The general pages describing the company, its products, and contact information will be normal http. All pages and scripts used by the authors and customers will be accessed with https using the highest encryption level available, 128 bit if available to the country the user is coming from. This server will maintain the user database and provide the authentication when an authorized user logs in. The only service required on this server is the web server and its components and the sshd daemon.

Database Host:

The database host maintains the most valued asset that GIAC Enterprises possesses, its fortunes. Access to this host must be protected so along with the armoring that the other hosts have done, this host will be placed inside the corporate firewall. The firewall will ensure that the only external host that can access the database host is the web server via sql scripts. More details of this connection will be presented in the security policy section. This host will also be accessed by internal corporate users in order to maintain the fortune databases and to keep track of what fortunes customers are retrieving and which suppliers are uploading new fortunes and business partner activities.

Mail Server:

The mail server that handles all of the corporate mail will be running the Microsoft Exchange server. This server will be behind the intranet firewall acting as a mail relay. The firewall will receive mail from the Internet and forward it to the internal mail server. The mail server will forward all non GIAC Enterprises mail to the firewall for forwarding to the Internet. Along with having a hardened OS, all patches installed, the mail host will also run Norton Antivirus software. This is a different antiviral software package than is run on the intranet firewall that provides additional virus detection. It is not unusual for a virus to be caught by McAfee but not by Norton and vice versus. Running both of these antiviral programs increases the chances of catching any viruses.

Syslog Host:

There will be a syslog host on each of the networks allowing the systems on those networks a centralized place for logging significant system events. The systems utilizing the syslog hosts will be the servers and routers on the subnet. The syslog hosts will not only be hardened like the other servers/hosts in the system they will also be hardened like the DNS servers and only have syslogd running. This will prevent any connections normally open when inetd is running. Periodically the logs of each of the syslog hosts will be moved and aggregated to a central syslog host on the GIAC

intranet. These transfers will be accomplished using the secure ftp portion of the open-ssh package. The configuration parameters for the syslogd daemon are in the /etc/syslog.conf file. Each event that will be sent to the syslog server is defined in this file along with the action that needs to be performed. Annoying events will merely be logged while severe events will be logged as well have alerts sent to key personnel via pagers. The termination of the httpd process on a web server would fall into the severe class of event.

IDS Sensors:

To provide feedback about unexpected activity on the networks and to provide information about the effectiveness of the layers of network security armored and dedicated Intrusion Detection System sensors will be positioned at strategic locations inside the corporate intranet firewall, within the public subnet, between the border router and the firewalls, and outside of the border router. Having an IDS sensor outside the border router will provide a good picture about activity occurring outside of our border router that is being stopped by the router as well as activity that successfully penetrates the various layers of the perimeter and internal defenses. The interface of this sensor will be a tap on the T1 line between the ISP and the border router. The details of the IDS sensor operation is outside the scope of this architecture other than to note they exist and where they will be placed. The IDs sensors will have two interfaces, one for transferring data to the IDS console for analysis and one for a passive tap of the network. The IDS sensor will only accept encrypted connections from the IDS console on the interface with the IP. The console will reside on the corporate intranet and will periodically aggregate the data collected by the sensors for the analysts to review.

Network Security:

The strategy utilized for network security will be a continued defense in depth and a choke point strategy. The choke point strategy limits the entry points of the networks to a limited number of points. By taking this strategy it is easier to monitor activities that are occurring where the corporate networks meet the Internet. The implementation of this strategy will mean that GIAC Enterprises will only use one T1 connection to its ISP for its Internet connection. The components of the network security will be the routers and switches used for connecting the subnets and the firewalls. To describe the network security I will start at the Internet and work my way into the internal corporate networks. This description will only cover the routers, switches, and firewalls it will not cover the hosts since they were covered in the previous section.

Border Router:

The border router that provides the connection to the ISP and the Internet will be a Cisco 3640. The Cisco 3640 router is a mid priced router easily capable of providing the throughput required for a T1 connection. The T1 interface module provides the CSU/DSU built in so there is no need for additional CSU/DSU equipment. This router also supports 10/100 auto sensing Ethernet connections to provide the required LAN speeds without collisions. When the traffic to the GIAC Enterprises increases and business expands this router is easily upgraded to provide up to 3 T1 interfaces. Since

this is the first line of defense from Internet attacks it is important that its configuration is hardened. To provide this hardening all unnecessary services will be disabled and although inconvenient all maintenance of the router will only be performed via the console port. Tftp could be used to configure the router, however I consider this protocol too risky and the process of configuring tftp on a host then unconfiguring to be prone to mistakes that will leave a large vulnerability in the system. This router will be configured for both ingress as well as egress filtering. The egress filtering will prevent any back doors that have been installed during a compromise from launching a spoofed IP attack from within the network. The border router will be connected to the border switch used to connect the firewalls that protect the other networks.

Switches:

All switches within the system will be Cisco 2900 XL switches. These switches provide the throughput required and the 10/100 Ethernet auto sensing feature. These switches also provide a SPAN feature that allows traffic on the ports of the switch to also be sent to a spanning port for monitoring. This feature will be used for purposes of sending traffic to the IDS sensors. Using switches for connecting the components of a subnet helps to limit the amount of information available to someone who has set up a sniffer on the network.

Firewalls:

The firewalls used in this architecture will be Gauntlet Version 5.5 proxy firewalls from Network Associates running on Sun Solaris 2.6. The hardware will be Sun Rack mounted Ultra 220s with dual processors. Initially there will be two firewalls; one firewall will sit between the border switch and the public servers, the public firewall, while the other firewall will sit between the border switch and the corporate intranet switch, the intranet firewall. The intranet firewall will be the mail relay host for the GIAC Enterprises mail domain. The Gauntlet firewall software uses a combination of smap, smapd, and sendmail, in send only mode, for relaying mail. The smap program listens on port 25 (SMTP), when it receives a message smap stores it in a uniquely named file. Periodically smapd runs, reads all of the files which have been created by smap and using sendmail set for send only delivers the mail. All of the programs used for handling the mail are running as uucp, a user with very limited privileges. This firewall also comes with built-in antiviral software using the Olympus engine for scanning and eliminating viruses from email.

VPN Device:

The user group identified as the newly acquired company, Joes Fortunes, resides in a different geographic location. This was a friendly merger and to support the process of merging the operations of Joes Fortunes with the operations of GIAC Enterprises the users on Joes Fortunes intranets will be as trusted as the users on the GIAC Enterprises intranets. To provide this seamless interface a VPN will be established between the intranet of GIAC Enterprises and the intranet of Joes Fortunes. The VPN will be created using BorderGuard 3000 routers from Blue Ridge Networks. These devices provide the VPN via encrypted tunnels between the intranets using ESP and will use triple DES for the encryption of the tunnel. The BorderGuard will sit between

the border switch and the intranet switch in parallel with the intranet firewall.

© SANS Institute 2000 - 2005, Author retains full rights.

Public Syslogd	Public syslog host	204.37.152.13
Public IDS	Public IDS sensor	204.37.152.6
Intranet Firewall	Firewall for corporate intranet	204.37.152.17 Outside IF 204.37.152.33 Inside IF
BorderGuard	VPN device at GIAC Enterprises	204.37.152.19 Outside IF 204.37.152.34 Inside IF
Intranet IDS	Intranet Ids sensor	204.37.152.35
Private Data Base	Corporate Data Base server	204.37.152.40
Mail Server	Corporate Mail server	204.37.152.42
Private DNS	Corporate DNS server	204.37.152.41
Intranet Syslog	Intranet syslog host	204.37.152.38
IDS Console	IDS analysis console	204.37.152.45
External syslog	External syslog host	204.37.152.21
External IDS	IDS sensor outside of any firewall	204.37.152.20
IDS Console	IDS analysis console	204.37.152.45
Joes BorderGuard	VPN Device at Joes Fortunes	204.37.154.19 Outside IF 204.37.154.34 Inside IF

The netmask used for the subnets involved for perimeter defense will be a 28 bit mask, 255.255.255.240. This provides each subnet with 14 IPs for hosts, one for the network, and one for broadcast. The ranges and their associated subnet is provided in the table below. The internal GIAC Enterprises workstations will use private IP ranges in the 172.21.0.0 range while Joes Fortunes will be using the 172.22.0.0 range. The intranet firewall will be using NAT for the users accessing sites on the Internet.

IP Range	Subnet Name
0-15	Public subnet
16-31	External subnet
32-47	Intranet subnet
48-255	Corporate internal use

Questions #2: Security Policy:

Introduction:

The security policy for GIAC Enterprises will be very restrictive for incoming traffic and moderately restrictive for outbound traffic. The strategies used will be that of least privileged and fail-safe. The concept of least privilege is based on the idea that a user only needs the minimal set of resources or permissions required to do their job or perform required functions within the system. Adding the strategy that if a failure occurs access will be denied, the fail-safe strategy [1]. The policies defined in this section are based on the user group descriptions defined in the Security Architecture. If additional information or explanation is required I will add it in this section, otherwise I

am included those descriptions by reference.

Policy Implementation:

Border router:

The border router is a Cisco router running the current IOS, Version 12.1(1)e, with all of the current patches applied. The configuration of the router will be accomplished by first installing a minimal configuration using the startup install provided by Cisco and then using the console port loading the actual configuration from a file which was created offline. The configuration of the router requires global configuration commands that apply to the router in general, interface definition commands, and Access Control List definitions that are tied to an interface by the interface definition. As mentioned before the configuration and administration of the router is performed from the console port. To accomplish this we will use a Sun workstation in the same area as the router. The Sun workstation will connect to the console port using the serial port TTYA. The tip utility will be used as the interface program. The tip parameters are specified in the /etc/remote file as follows:

```
Border-rtr:\
      :dv=/dev/term/a:br#9600:el^C^S^Q^U^D:ie=%$:oe=^D:ec:
```

where :dv specifies the device

:br is the speed

and :ec specifies to initialize echo check so tip synchronizes with the remote host during transfers. This is necessary so the tip program will wait for the router to finish processing the commands before sending the next command. If this is not in the remote configuration the tip session will overrun the buffer on the router and the configuration file will not be successfully transferred.

The global configuration commands will prevent the router from using services that are known to be vulnerable to attack or ways used to attack servers behind the router or which are rarely used. The services that are rarely used may have future vulnerabilities. These services include snmp, httpd and bootp on the router, echo, discard, chargen, and daytime. It is also important to prevent a smurf attack. Sending directed broadcasts to servers on a network which then propagate these broadcasts performs a denial of service called a smurf attack. Denying ip directed-broadcast on the router can prevent this attack. Denying ip unreachable will prevent releasing network information from the router. The next global configuration we need is a banner that clearly states that only authorized personnel are permitted access to the router. The final global setting required is the syslog setup. The setup of the logging facility defines the log host, the log facility, and the level of logging to be done. The Cisco site <http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/as5xipmo/sysmgt.htm> can provide more details on setting up the management protocols for a Cisco router. The Solaris man pages for syslog.conf will also provide details about setting up the syslog facility on the syslog server. Along with all of the denying we do need to allow ip routing or packets won't go anywhere. Along with the IP routing being allowed we also need to provide the routes needed to deliver our packets, the static route to the public

subnet, this will be to the public firewall which will then route the packets to the servers, and the default route. The intranet firewall will do a NAT for any connection originating from within the intranet so no routes are required for the intranet since the destination IP will be the intranet firewall's outside IP. The following is the set of configuration commands that will limit the services mentioned above:

```
! This is a comment line within the router configuration file.
! Setup the host name
hostname giac-rtr
! Prevent password display in the clear when the configuration is
! displayed on the console.
Service password-encryption
! Enable a separate password for privileged commands
Enable secret
! Deny rarely used services. These are turned off by default in IOS 12
no service tcp-small-servers
no service udp-small servers
! Limit ICMP responses
no ip unreachable
! Deny other vulnerable services
no bootp
no httpd
no ip source-route
no ip finger
! Do not run Cisco discovery protocol
no cdp
! Setup logging facility Default is to log to facility local7
logging 204.37.152.21
! Define the banner to display
banner / WARNING: Only AUTHORIZED users are permitted system Access. /
! Routing. Need to route public subnet to the outside interfac of the public
firewall
! Note: The intranet firewall uses NAT with the outside IP on the subnet
with ! one of the router interfaces so the intranet IP's do not need a
route.
! The syntax is "ip route destination subnet-netmask gateway-IP
ip route 204.37.152.0 255.255.255.240 204.37.152.18
! Provide a default route to the ISP via the T1 serial connection
ip route 0.0.0.0 0.0.0.0 serial0/0
! Setup the console access and deny all other administrative access
!Define console behavior
line con 0
login
password a;dfhoiehfhelioqwe
exec-timeout 5 0
!
!Disable logins on telnet ports
```

```

line aux 0
login
access-class 10 in
!
line vty 0 4
login
access-class 10 in
!
line vty 5 15
login
access-class 10 in
!

```

Now that the general configuration is completed we can setup the interfaces. Each interface will need to be defined based on its interface type. Each interface gets an assigned IP along with the associated ACL list that has been defined for it. The ACL will use the in option so the access controls will be performed prior to the packet being routed. This will help reduce the load on the router's processor.

```

! Setup the Ethernet interface
interface FastEthernet 1/0
no shutdown
description connected to Border_Switch
ip address 204.37.152.30 255.255.255.240
ip broadcast -address 204.37.152.31
! Deny Smurf attack
no ip directed-broadcast
keepalive 10
! Associate the access list with this interface
ip access-list 102 in
! Setup the PPP connection and the serial T1 interface
interface Serial 0/0
no shutdown
description connected to Internet
service-module t1 clock source line
service-module t1 data-coding normal
service-module t1 remote-loopback full
service-module t1 framing esf
service-module t1 linecode b8zs
service-module t1 lbo none
service-module t1 remote-alarm-enable
ip address 204.27.68.253 255.255.255.252
encapsulation ppp
! Associate the access list with this interface
ip access-list 101 in

```

We will now setup the ACLs specified above. There are two types of ACLs on a Cisco router, standard and extended. Both ACLs define what traffic is permitted or denied access. The difference between the two types is that the Standard type are numbered from 1 to 99 and allow the definition of the source and destination IPs but do not define the protocol or ports allowed. The extended access lists are numbered from 100 to 199 and define the source and destination IPS as well as the protocol and port to be permitted or denied. Since the security policy is to be very restrictive we will use the Extended ACL format. This will allow us to not only restrict IP access at the router but also the ports. The default for a Cisco router is that if there is an ACL defined all traffic is denied unless specifically permitted. This provides part of the fail-safe strategy we are using. The order of the ACLs is also important. The IOS scans the ACLs and uses the first match to decide access; it also adds all new ACL entries at the bottom of the list. If an ACL to allow access needs to be added it is best to place the ACL in the proper position within the configuration file and then load the whole file at once via the console port. It is also advisable to create a separate file for the main configuration that includes the interface and global definitions and a file for each ACL. Having these separate files will reduce the load time when an ACL is changed and also keep the file size small so it will be easier to review each ACL. The approach of the ACL is to prevent unexpected traffic from entering or exiting our networks. The traffic that would not be expected to enter our networks are the IPs that are part of our network. If we see a packet with a source IP that is within our network IP range we know it is a spoofed packet and need to reject it. We also know that if we see a packet with a source IP that is not within our IP range exiting our networks, this is also a spoofed packet and should be stopped. By preventing these spoofed packets from exiting our networks, will help prevent a denial of service attack using spoofed IPs from originating from our network. Packets originating from non-routable IP ranges also need to be stopped. When creating the ACL, it is important to remember that the netmask is opposite of what most think of when they specify the netmask for routing, a 0.0.0.0 netmask exactly matches the IP while a netmask of 255.255.255.0 matches all IPs for that subnet. The following is the format of an Extended ACL:

```
Access-list <access-list number> [permit|deny] <protocol> <source ip> <source netmask> <destination ip> <destination netmask> [eq|gt|lt] <port>
```

A Standard ACL has the following format:

```
Access-list <access-list number> [permit|deny] <source ip> <source netmask> <destination ip> <destination netmask>
```

The ACLs we will use follow:

```
! Access List for traffic coming from the Internet
! First we remove the existing access list to ensure the list is what we expect
no access-list 101
! Deny any attempts at spoofing and non-routables.
! log these attempts since they are suspicious
access-list 101 deny 192.168.0.0 0.0.255.255 log
access-list 101 deny 172.16.0.0 0.15.255.255 log
access-list 101 deny 10.0.0.0 0.255.255.255 log
access-list 101 deny 127.0.0.1 255.255.255.255 log
```

```

access-list 101 deny 204.37.152.0 255.255.255.0 log
!
!   Establish Ingress filtering
!   Only permit expected ports from the outside to our known IPs
!
!   Allow the VPN connection between Joes BorderGuard and GIAC
BorderGuard
access-list 101 permit 50 204.37.154.19 0.0.0.0 204.37.152.19 0.0.0.0
!   Access to the DNS server Connect for DNS queries
access-list 101 permit udp 0.0.0.0 255.255.255.255 204.37.152.10 0.0.0.0 eq 53
!   Access to the public Web server Connect for http
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.11 0.0.0.0 eq 80
!   Now allow high number followup ports but deny X11 ports
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.11 0.0.0.0 gt
1023
access-list 101 deny tcp 0.0.0.0 255.255.255.255 204.37.152.11 0.0.0.0 gt 6000
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.11 0.0.0.0 gt
6099
!   Access to the business Web server Connect for https
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 eq
443
!
!   Now allow high number followup ports but deny X11 ports
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt 1024
access-list 101 deny tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt 6000
access-list 101 permit tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt
6099
!
!   Allow inbound access to the Intranet Firewall
!   The only inbound connections we allow to the intranet will be the
!   return packets from connections which originated from the intranet
!   udp domain and mail to the GIAC
!   Enterprises mail domain.
!   Allow mail to be sent in SMTP
access-list permit tcp 0.0.0.0 255.255.255.255 204.37.152.17 0.0.0.0 eq 25
access-list permit udp 0.0.0.0 255.255.255.255 204.37.152.17 0.0.0.0 eq 53
!   Now allow high number followup ports but deny X11 ports
access-list permit tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt 1023
access-list deny tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt 6000
access-list permit tcp 0.0.0.0 255.255.255.255 204.37.152.12 0.0.0.0 gt 6099

!
!   Perform Egress filtering
!   Allow the BorderGuard VPN out
access-list 102 permit 50 204.37.152.19 0.0.0.0 204.37.154.19 0.0.0.0

```



```

!   Since the Intranet Firewall does NATing we only need to allow the outside
!   IP of the firewall out. We will allow tcp Internet access
!   and udp for DNS since the Firewall is acting as a DNS forwarder for the
!   Intranet DNS server.
!   The Public firewall does not do NATing
!   so we need to let all of the web server IPs out
!   and we need to allow udp out for the DNS server queries
!
access-list 102 permit tcp 204.37.152.17 0.0.0.0 0.0.0.0 255.255.255.255
access-list 102 permit udp 204.37.152.17 0.0.0.0 0.0.0.0 255.255.255.255
access-list 102 permit tcp 204.37.152.11 0.0.0.0 0.0.0.0 255.255.255.255
access-list 102 permit tcp 204.37.152.12 0.0.0.0 0.0.0.0 255.255.255.255
access-list 102 permit udp 204.37.152.10 0.0.0.0 0.0.0.0 255.255.255.255 eq 53
!
!   explicit deny any outbound traffic not permitted above. Log this so we can see
!   if one of the internal hosts may have been compromised
access-list 102 deny ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 log

```

As can be seen by the above ACLs we are very restrictive at the border router. Some will say that this is not needed since the firewalls will also be applying these restrictions. By adding these restrictions at the border router we reduce the network load on the firewall and apply the rule of redundancy so if one filter does fail and allow something through there are more filters behind. We also prevent some of the standard Ddos attacks on the firewall. The denying of ICMP traffic at the router, although making troubleshooting more difficult does help remove some of this annoying traffic.

Firewall Configuration:

The firewalls used for this application are Network Associates (NAI) Gauntlet Version 5.5 firewalls. Gauntlet is a proxy firewall so any connections that are allowed will be handled by the firewall making the connection to the requesting client and then creating a separate connection to the server thus the client never actually has a connection directly to the server. The proxy then passes the communication streams between these two connections. They are installed on the Sun Solaris 2.6 OS using the installations provided by NAI. They require that Sun Solaris be installed selecting the developer option for the level of installation. A complete installation of Solaris including OEM support is not performed. Following the installation of Solaris the current set of patches is installed. These patches include all of the recommended patches and the current set of security patches. The next step is to install the Gauntlet software. This software is installed from the provided CD. All of the Gauntlet management and configuration files are installed in the /usr/local/etc directory. The installation also creates startup files in the /etc/rc2.d directory. Startup files which are not to be started are moved by the installation script to files named disabled.XXX, where XXX is the original name of the file. Once the install is complete the machine is rebooted. After rebooting and logging in as root Gauntlet automatically starts its initial

configuration program. This program allows you to specify the interfaces you are going to use for the inside, the outside, and the firewall manager user ids and passwords along with the permitted IPs for the management gui used to configure the rule sets for the firewall. You can also select the proxies that will be enabled and the trusted networks. However configuring proxies and the trusted networks is easier using the epsm-gui program. This program provides a graphical interface for configuring the firewall, figure 2.0 show the login screen for the epsm-gui program. There is a Windows version as well as a Solaris version. The program authenticates each connection using a userid/password as well as the source IP with all data exchanges being encrypted.

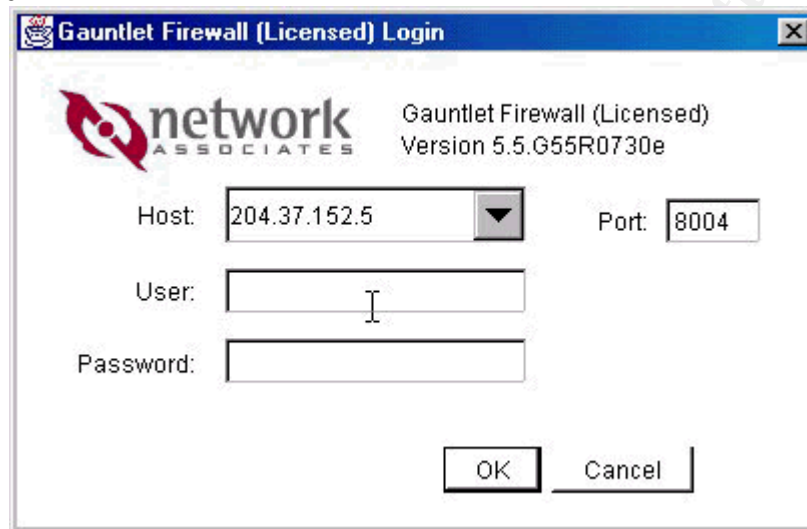


Figure 2.0

Once the basic install is complete the firewall is once again rebooted. During this reboot Gauntlet applies its default filter set which is to deny everything and starts any proxy programs which had been enabled. Now that the initial setup of the firewall has been completed and the firewall rebooted, the configuration of the firewall will be done using the epsm-gui program. The configuration of Gauntlet is based on object concepts with a set of objects being combined to determine access to the firewall proxies. The object groups are rules, rule elements, and services. The rules contain the source and destinations rules, while the rule elements contains the networks, network groups, and service groups. There is a screen for configuring each of these object groups. The screen for the services provides the interface for enabling or disabling the supported services, telnet, http, ftp etc. The screens for each service also provide for defining the logging to be performed for that service and parameters for timeout of an inactive connection. Figure 2.1 shows the interface for selecting the services. The enabled services are in green. This does not indicate if the service is outbound or inbound just that they are enabled. The parameter for number of concurrent connections for that service is also specified here so performance tuning can be tweaked.

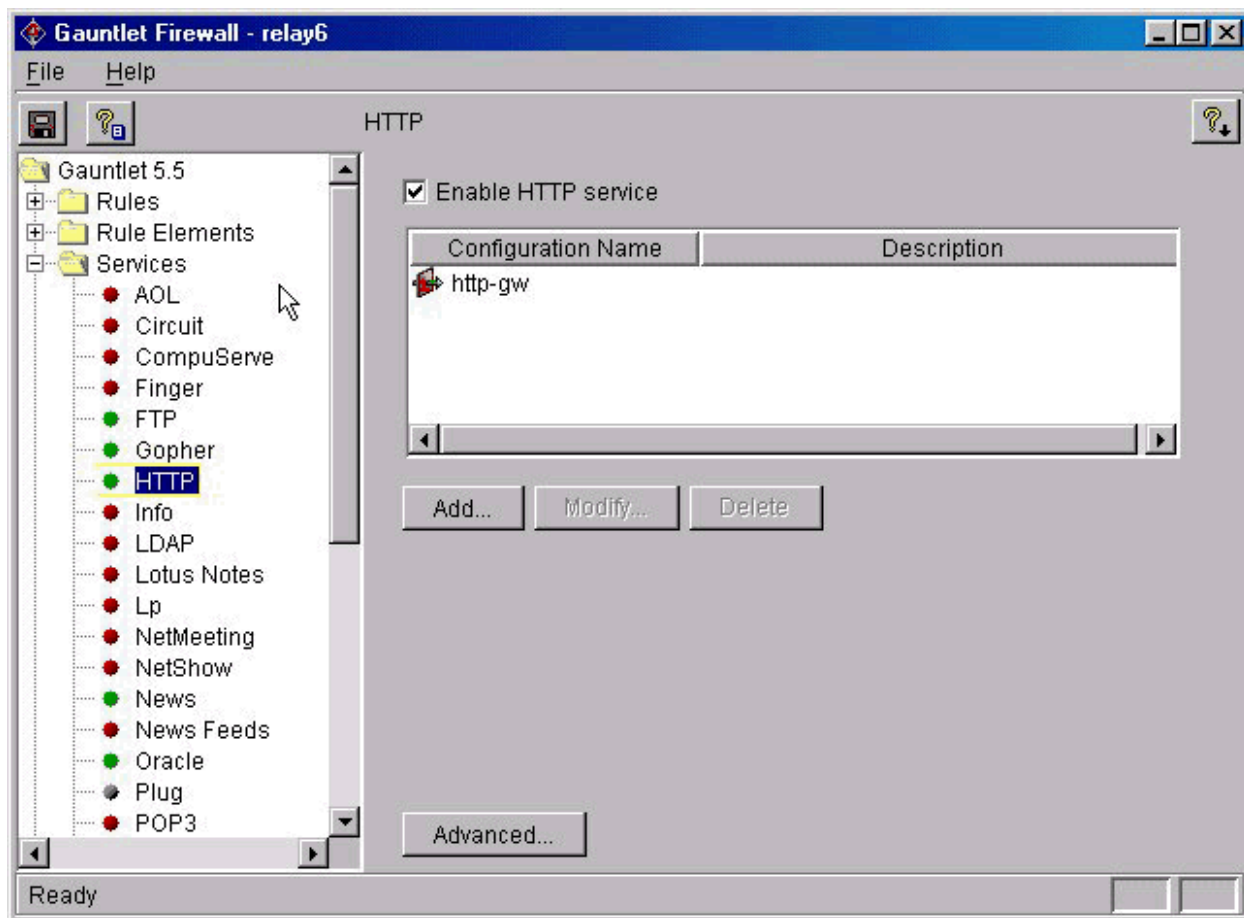


Figure 2.1

The networks object defines the networks to be used within the rule sets. Networks are aggregated into a network groups object, The initial network groups are trusted and untrusted. Additional networks groups can be defined if needed. This set of objects defines a network using the IP and wild cards along with the netmask. When a network is defined it is placed into a network group, trusted or untrusted for example. By default the inside subnet is in the trusted group while all other networks are in the untrusted group.

Service groups are created to define what might be considered user types. The service group defines the services that are permitted for a particular user group. The display for defining and modifying service groups specifies the available services on one side with the permitted services on the other side. Services can easily be moved back and forth using arrow buttons between the two windows, see figure 2.2 and figure 2.3.

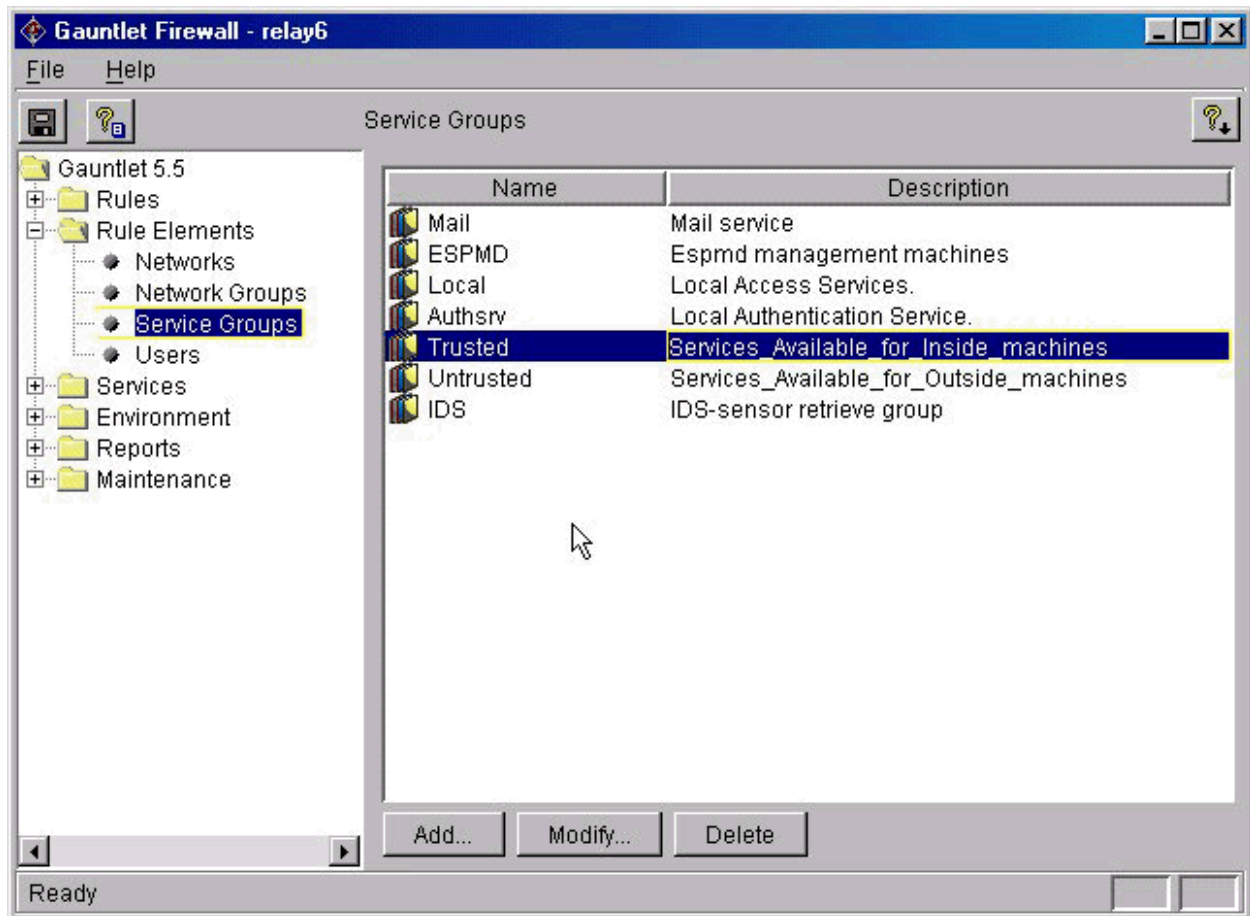


Figure 2.2

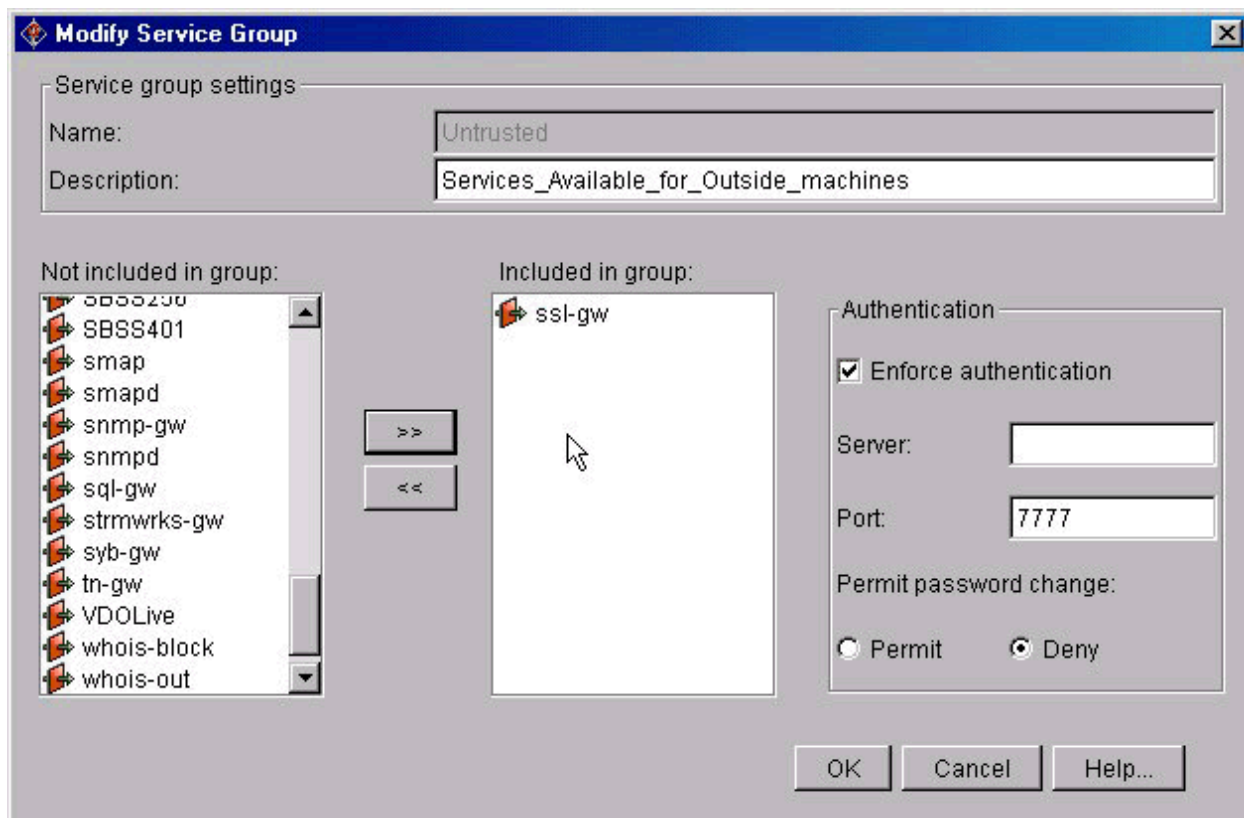


Figure 2.3

The source rules connect the network groups and links them to the service group. By creating this link Gauntlet can examine the source IP of the packet and determine if the requested service is allowed. The source rule will specify that the trusted networks be permitted to use the services within the trusted service group, figure 2.4. So before you can create a source rule the network group and service groups must be defined. To provide a fair amount of granularity individual IPs and services can be permitted or denied in the source rules.

© SANS Institute

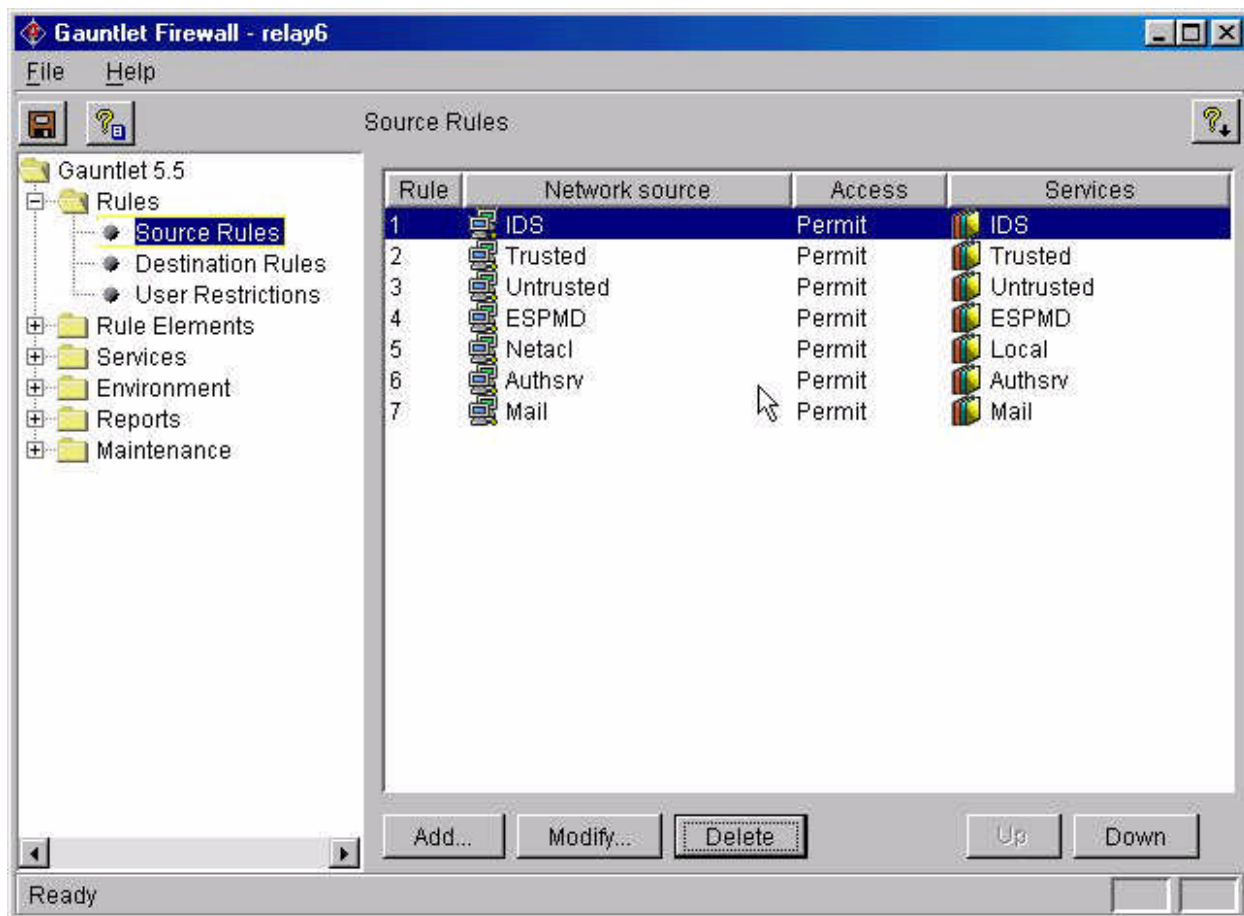


Figure 2.4

The final part of the access controls on a Gauntlet firewall is the set of destination rules figure 2.5. These rules define the destination the various network groups are permitted to access. Like the source rules the network groups need to be defined prior to creating the destination rules. An example of a destination rule would be that the trusted network group is permitted to access all networks. Wildcards are permitted in the definition of the destination IPs.

© SANS Institute 2000

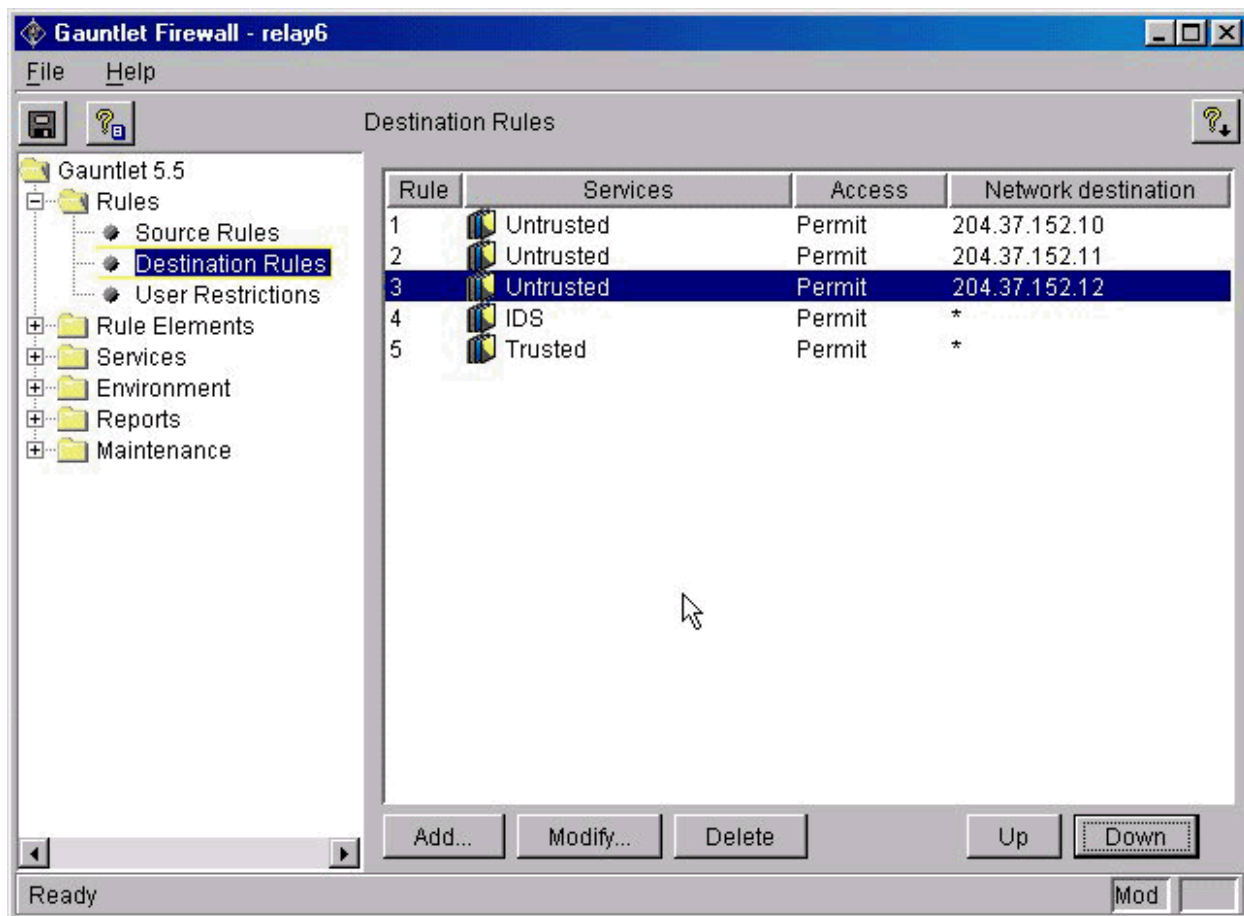


Figure 2.5

The order of the rules in the source and destination rules is significant. The first match for a rule will be the rule selected so it is recommended that the more specific rules be placed first with the more general rules last. Since the network IP drives the rules, the more specific networks need to be first in the list. The gui interface provides places for the entry of the fields relevant for each rule. For the source rules there are drop down menus with the various network groups and service groups or a place to select and enter a specific IP or network. The services selection contains radio buttons for selecting whether to permit or deny access and a drop down menu to select the service group or specific service. Both of the firewalls will be configured as DNS forwarders.

Now that the general setup of the Gauntlet firewall has been covered the specifics for the Public and Intranet firewalls will be addressed.

Public Firewall:

The Public firewall will need three service groups to implement the security policy; trusted, untrusted, and administrators. The trusted group will be for connections originating from the public subnet outbound. These connections will be from the users that are maintaining the public servers and will be logged in locally to the servers. These users will be permitted http/https, ftp, telnet sshd, and sqlnet for connection to

the data base server. The untrusted service group will be the users originating from the Internet or users on the GIAC intranet browsing the web sites on the public servers. The untrusted service group will only have access to http/https and SSL services. The final service group is the administrator service group. This group will administer the public servers remotely. The access and only service available to this group is the ssh service. Since Gauntlet does not have an actual ssh proxy a plug will be defined for port 22 to allow this service. The definition of the plug specifies the port and source and destination IP. We will specify the intranet firewall outside IP as the source since all administration will originate from within the GIAC intranet and use a wildcard to specify all of the public subnet.

Corresponding to the service groups there will be three source rules. All of the rule definitions and group creation/definitions were done using the espm-gui interface. After all of the rules and objects have been defined the rules are applied to the firewall configuration by selecting the save and apply option from the file menu. If you exit the espm-gui without saving you will be prompted if you want to save and apply the configuration. The following tables summarize the groups and rules defined:

Service Groups:

Service Group	Services
Trusted	http, ftp, telnet, SSL, sqlnet
Administrators	ssh
Untrusted	http/https SSL DNS

Network Groups:

Network Group	IPs
Administrators	204.37.152.17
Trusted	204.37.152.0/28
Untrusted	*

Source Rules:

Network Group	Permit/deny	Service group
Administrators	Permit	Administrators
Trusted	Permit	Trusted
Untrusted	Permit	untrusted

Destination Rules:

Network Group	Destination
Administrators	204.37.152.0/28
Trusted	*
Untrusted	204.37.152.10
Untrusted	204.37.152.11
Untrusted	204.37.152.12

Intranet Firewall:

The intranet firewall is more restrictive for inbound traffic than the public firewall. The

only inbound connection that will be allowed will be an sqlnet connection originating from the business web server 204.37.152.12 with a destination of 204.37.152.40. To make this connection the Business server will request a sqlnet connection using the outside IP of the intranet firewall, 204.37.152.17. The intranet firewall will be listening for this connection on its outside interface, when the connection request comes the configuration of the sqlnet proxy identifies the data base server along with the data base ID on that server. The access rules are checked and if the requestor is permitted the sqlnet proxy will complete the connection to the data base server. The outbound traffic will permit the same services that are permitted from the inside of the public subnet. Guantlet requires that at least a trusted and untrusted network group and service group be defined. Like the public firewall ssh will be defined using a plug proxy. Since the ssh plug will be outbound it will allow connection to all hosts/IPs and must originate from the GIAC intranet. Defining the trusted network requires that we define three subnets in order to permit all of the 204.37.152.0 IP range excluding the two subnets, 204.37.152.0/28 and 204.37.152.16/28 which are outside of the intranet and are considered as untrusted relative to the intranet firewall.

Service Group	Services
Trusted	http, ftp, telnet, SSL, sqlnet, ssh
Data Base	sqlnet
Untrusted	(none)

Network Groups:

Network Group	IPs
Trusted	204.37.152.128/25 204.37.152.64/26 204.37.152.32/27
Data Base	204.37.152.12
Untrusted	* 204.37.152.0/28 204.37.152.16/28

Source Rules:

Network Group	Permit/deny	Service group
Data Base	Permit	Data Base
Trusted	Permit	Trusted
Untrusted	deny	untrusted

Destination Rules:

Network Group	Destination
Data Base	204.37.152.40
Trusted	*

BorderGuard:

The VPN device that will be used for creating the VPN tunnel between the GIAC Enterprises intranet and the Joes Fortunes intranet will be a Blue Ridge Networks BorderGuard 3000 located in parallel with the intranet firewall at each location. The BorderGuard has two Ethernet interfaces one connected to the intranet switch and the other connected to the external switch. To establish the VPN using the BorderGuards requires four steps:

1. The BorderGuards are configured
2. The BorderGuards generate and exchange keys
3. The routing on the intranets is setup to direct traffic destined to the intranet at the other end of the VPN to the local BorderGuard
4. The BorderGuards are deployed and when a packet destined for the distant end of the VPN is received, the BorderGuard establishes the encrypted tunnel and delivers the packet. This tunnel will stay established for as long as traffic is crossing the tunnel or until the timeout parameter is reached due to inactivity.

Configuring the BorderGuard:

The BorderGuard comes with a quick start script that runs when a new Borderguard is turned on. This script provides the basic configuration of the device internal and external ip netmask hostname etc. When a configured BorderGuard starts it looks for a number of files on its hard drive. These files provide the configuration it uses for providing the VPN connections. The startup file controls the startup and initialization of the BorderGuard. This is the first file read and the commands within it are executed when the Borderguard starts. The BorderGuard has four command categories, system, net sentry (ns), ip, and data privacy facility (dpf). The net sentry commands define the filters used on packets entering the BorderGuard while the dpf commands define the VPN tunnels. Notice that during the initialization a filter is set to drop all packets coming into the device before the interfaces are started. This ensures that it is very difficult to exploit the device between when the interfaces are started and the real filters are put in place. The following are the files that provide this information. A “\” at the end of the line is a continuation character.

Startup:

The initialization file.

```
# general initalizations
#
system set name "giac-vpn.giac.com"
system set default severity 8
system set default facility all
system set con sev 8
#
# compile filters and apply the filter to drop all packets coming to the box.
#
ns compile filters.ip
ns apply ip forgetit on last
#
```

```
# start interfaces
#
ip start if en01 204.37.152.37 netmask 255.255.255.240
ip start if en02 204.37.152.19 netmask 255.255.255.240
#
# STATIC DEFAULT ROUTE to external giac-rtr
ip define route 0.0.0.0 204.37.152.30
ip set ipsendredirects off
ip set ipreceiverredirects off
ip set ipsendudpdiscards off
system set default more off
#
# NETSENTRY
@apply.ip
```

```
dpf establish sleeve giacbg2joesbg
```

Apply.ip:

This file contains the net sentry filter commands. It is executed as a result of the @apply.ip command in the startup file. The commands in this file apply a filter based on the ip addresses in the packet. The syntax is:

Apply ip <filter-name> on <subnet or IP> netmask <netmask> [to|from|tofrom] <subnet> netmask <netmask>

The on subnet refers to the subnet of the interface the packet was received from and the to|from|tofrom refer to the direction the packet is headed. Apply the filter when it is coming from someplace, from, or if the filter is going to someplace, to. The last set of IP and netmasks provide the to or from information.

```
#
netsentry compile filters.ip
netsentry unapply ip all
netsentry apply ip TOJoes on 204.37.152.32 netmask 255.255.255.240 \
    to 204.37.154.32 netmask 255.255.255.240
netsentry apply ip FMJoes on 204.37.152.32 netmask 255.255.255.240 \
    from 204.37.154.32 netmask 255.255.255.240
#
# Allow in anything from the sleeve
netsentry apply ip FMJoes on *.*.* from 204.37.152.19
# then router destined traffic to/from the world
netsentry apply ip rtrTraffic on 204.37.152.19 tofrom *.*.*
netsentry apply ip rtrTraffic on 204.37.152.37 tofrom *.*.*
```

Sleeves:

The sleeves file provides the definitions for how the encrypted tunnel is configured. The

sleeve default command sets up the default parameters. The key parameters here are the timeouts, in minutes, and the encryption of the tunnel using triple_des_cbc. It is important that the sleeve definition line is the same in the borderGuards at both ends. The sleeve is established at the end of the startup file. The sleeve name must match. The syntax is:

```
Sleeve <sleeve name> group1 <outside IP > group2 <outside IP>
#
sleeve default protocol ip key short idle_timeout 10 \
    key_lifetime 1440 status_timeout 60 connect_timeout 60 \
    replay_prevention none encryption triple_des_cbc compression none \
    integrity none

#
sleeve giacbg2joesbg group1 204.37.152.19 group2 204.37.154.19
```

Filters:

The last file configured is the filters file. This file provides the filter definitions for the filters that were applied in the apply.ip file. This file has the filters for both borderguards. The filters file can be the same but the apply.ip file will be unique at each site. You just have to be careful to get the right filter applied to the right interface relative to the borderguard you are configuring.

```
#
# filter forgetit: filter will alarm and fail all packets
#
filter forgetit
    alarm 9 fail;
end

# filter for incoming traffic
filter onsleeve
#    Allow only encrypted traffic
    not ip_protocol in (50) alarm 10 fail;
#    Do not allow source routed packets
    ip_option_present 0x89 alarm 10 fail;
    ip_option_present 0x83 alarm 10 fail;
end
#
# General explanation of syntax:
# if packet is going to GIAC put the packet on the sleeve.
# This filter would be applied at the borderguard at Joes Fortunes.
filter TOGiac
    any set_sleeve giacbg2joesbg break;
end
#
# If this packet has a GIAC source IP but did not come from the sleeve discard
it.
```

```
# It is probably spoofed.
filter FMGIAC
    not sleeve giacbg2joesbg alarm 10 fail;
end
# if packet is going to Joes Fortunes put the packet on the sleeve.
# This filter would be applied at the borderguard at GIAC Enterprises..
filter TOJoes
    any set_sleeve giacbg2joesbg break;
end
filter FMJoes
    not sleeve giacbg2joesbg alarm 10 fail;
end
```

Generating and exchanging keys:

The generation and exchange of keys is straightforward and can be done in a lab where the two borderguards can be connected to a network without subjecting them to the Internet. Once the borderguards are configured and booted enter the command:

```
Dpf generate keys
```

This will generate the public and private keys for the encryption process. The keys are based on the IP so if you change the IP of the borderguard you will need to regenerate and exchange the keys again.

To exchange the keys the following commands are entered.

```
Dpf introduce mykeys to 204.37.154.19
```

This will present your keys to the borderguard with the ip specified. When the borderguard receives this message it will display a message on the console that it has received introductions from 204.37.152.19. At this time if you are expecting introductions you type the command:

```
Dpf receive introductions
```

The borderguard will accept the introductions and load the public key into its active public key database. This command must be done at both of the borderguards. Once the introductions are received the keys have been exchanged. It is important to unload the keys at this time. Unloading the keys will store them on disk so if the borderguard is powered off, when it is restarted it will have the keys for performing the encryption. The outside IP is used for this process.

Establishing the routing on the GIAC intranet is beyond the scope of this paper and we will just say that some magic is done by the network team and the routing is correct.

Once the borderguards have been configured and deployed you should be able to ping a host at the distant end. Since pings are not allowed through the firewall or the giac-rt if it is successful we can state that the VPN is working.

Host Details:

In addition to the hardening of the OS the userid used for running the web server will be web. This is a user with little or no privileges within the system other than to access the files needed by the web server process. By using this userid if the web server process

is compromised the intruder will have limited access to the system. The definition of the userid is usually in the config files for the web server or when the system boots the startup command for the web server can be done using the `exec su - web-userid` command in the startup script file.

The DNS server will be configured to not allow any zone transfers outside of the GIAC domain. This is done by adding the "allow-transfer {none}" to the configuration file `/etc/named.conf` on the DNS servers. Since both of the DNS servers are behind the firewalls which are acting as the DNS forwarders the configuration command to add to the `named.conf` file is "forwarders {204.37.152.19}" needs to be added to the intranet DNS while the public DNS needs is "forwarders {204.37.152.5}.

Question #3: Security Architecture Audit.

Introduction:

There are many reasons to perform a security audit; you want to determine if the current configuration is still the configuration which was originally loaded, you need to determine if new exploits have been discovered and could be used against the firewall, or the firewall has had a configuration change and you want to ensure that the changes have not introduced any weaknesses. For whatever reason security audits should be performed on a regular basis.

Approach:

The security audit will be performed over several days during normal work hours for interviewing staff and after hours on a weekend. The interviews with the normal duty staff is intended to determine what the security stance is expected to be and also to find out what the procedure is when an incident occurs. The scanning and testing of the systems against the security policy stated, what I will term the audit, should be conducted on a weekend or after hours for a couple of reasons. Conducting the audit on a weekend will minimize the impact if problems are encountered during the audit. These problems could be as simple as a slowdown in the response of the firewall to a complete failure of the firewall. If the firewall should fail the impact to the company is minimized and there will be time to repair the firewall or other equipment being tested during the audit. The second reason for performing the audit on the weekend is the response of the firewall will be faster and the tests and scans will take less time than during a workday when the firewall is loaded. The assessment will check for security holes in the configuration of the operating system, check the checksum database created when the system was first installed or the last database available, and port scans will be done to determine if there are listeners on unexpected ports.

All of the hosts need to have the following procedure used for determining if there are any vulnerabilities. I will go into more detail for the firewall since there are additional checks that must be performed. The initial audit performed will be a check of the firewalls OS to determine if there are any holes open in the configuration. The COPS and TIGER programs are good programs to check the security stance of the Solaris operating system. These programs will check many of the currently known

vulnerabilities of the OS as well as the permissions of files. Along with these programs sunsolve should be run. This program supplied by SUN will perform an audit of the patches that are installed and compare them with the available patches. This is a good way to ensure that the current set of patches has been installed. When sunsolve does its checks it also checks the packages that are installed so it is only checking the patches for the packages that are installed. Some caution needs to be exercised at this time however. When gauntlet is installed it replaces some of the low level drivers with hardened drivers it has developed. When installing the Sun patches you need to be careful not to replace these hardened drivers with a less secure one in a patch. NAI has information that specifies which of the Sun recommended patches are safe to apply. The version of the applications running on the servers needs to be checked and any application specific patches need to be checked. For the firewall Gauntlet needs to be checked as well along with a check of the Gauntlet patches that have been released. The gauntlet-version command will print out the version of each of the proxies in Gauntlet. In addition to the gauntlet-version command you need to cat or list the /usr/local/etc/mgmt/patchlog file. This file contains a list of all the gauntlet patches that have been applied and the date they were applied. If any required patches have not been installed they should be installed provided they do not replace a file that has been hardened. Also any issues uncovered by the running of COPS and TIGER should be addressed.

The next item to audit is the tripwire checksum database. This database will either be the database created when the firewall was originally created or one that was created the last time a change was made on the firewall. If there are any discrepancies they need to be researched and explained otherwise a compromise should be suspected and investigated.

Once the integrity of the operating system has been completed the rule sets should be examined to determine what they are supposed to be blocking and if the configuration appears correct. The next step is to run scans against the firewall from each of the network segments; the Internet, the public subnet, the external subnet, and the intranet. The results of the scans are compared to the expected list of ports expected to be open to determine if there are any differences. Any differences are required to be resolved. When a system is first deployed these same steps should be performed and the results from the scans on the original system should be compared. This resolution may be the creation of a new rule or the correction of any rule that was permitting access to the port. The resolution may also be the removal of a startup script which had been inadvertently added. Of course when an unexpected startup script appears, you need to determine where it came from. The log files also need to be checked to ensure that the correct level of logging was performed ie: alerts were logged for attempts to connect to an unauthorized port and the successful connections were logged. The port scans can be done using Hackershield or nmap. For this example I used nmap obtained from SunFreeware, <http://sunfreeware.com>, from a Sun Solaris machine. The results follow:

The -h option provides a quick reference of the options for nmap.

```
root (www) > ./nmap -h
```

```

Nmap V. 2.54BETA28 Usage: nmap [Scan Type(s)] [Options] <host or net
list>
Some Common Scan Types ('*' options require root privileges)
  -sT TCP connect() port scan (default)
* -sS TCP SYN stealth port scan (best all-around TCP scan)
* -sU UDP port scan
  -sP ping scan (Find any reachable machines)
* -sF,-sX,-sN Stealth FIN, Xmas, or Null scan (experts only)
  -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
  -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)
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

```

Running nmap and giving an IP or hostname to scan performs the simplest scan. This scan will check the most common ports. For completeness the scan needs to be done using the `-p` option to specify the ports to scan, `nmap -p 1-65356` will scan all the ports. The `-o` option can also be used so the output is saved to a file.

```

root(www)>nmap 204.37.152.18

Starting nmap V. 2.54BETA28 ( www.insecure.org/nmap/ )
Interesting ports on publicfw.giac.com (204.37.152.18):
(The 1534 ports scanned but not shown below are in state: closed)
Port      State      Service
21/tcp    open       ftp
22/tcp    open       ssh
23/tcp    open       telnet
25/tcp    open       smtp
53/tcp    open       domain
70/tcp    open       gopher
80/tcp    open       http
111/tcp   open       sunrpc
113/tcp   open       auth
389/tcp   open       ldap
443/tcp   open       https
444/tcp   open       snpp
6000/tcp  open       X11
32771/tcp open       sometimes-rpc5

```

It is also important to scan the UDP ports. The `sU` option is used to scan the UDP

ports. Once again all of the ports need to be scanned, although it will take quite some time for the scan to complete.

```
Nmap run completed -- 1 IP address (1 host up) scanned in 5 seconds
root(www)>nmap -sU 204.37.152.5
```

```
Starting nmap V. 2.54BETA28 ( www.insecure.org/nmap/ )
Interesting ports on publicfw.giac.com (204.37.152.18):
(The 1442 ports scanned but not shown below are in state: closed)
Port      State      Service
53/udp    open       domain
111/udp   open       sunrpc
123/udp   open       ntp
146/udp   open       iso-tp0
177/udp   open       xdmcp
421/udp   open       ariel2
514/udp   open       syslog
561/udp   open       monitor
1467/udp  open       csdmbase
3986/udp  open       mapper-ws_ethd
32773/udp open       sometimes-rpcl0
```

```
Nmap run completed -- 1 IP address (1 host up) scanned in 2544
seconds
root(www)>
```

For this example I only used the default and simplest form of nmap and found several items that need to be addressed. After running nmap the first time it was found that many ports were listening on the outside. The default for Gauntlet is for any enabled proxy to listen on all interfaces defined. It is possible to specify the interface for the proxy to bind with so in cases where access to that proxy would only be permitted on one of the available interfaces, the proxy should be bound to that interface. To provide an example I bound the FTP proxy to the inside (Trusted) interface and reran nmap. Now you can see that the FTP port is no longer open. Since the security policy only specifies that SSL and UDP DNS be allowed in from the outside all of the other proxies need to be bound to the inside interface. By doing this it will remove the processing required for the firewall to check the rules before denying access and logging the event. When the proxy is not listening on the interface the firewall will only need to log a connection attempt on an unserved port. All of the open UDP ports need to be examined and a determination as to why they are open needs to be made. We expected that UDP 53/domain would be open but the other ports were surprises especially the Sunrpc port 111/udp. This port certainly may be a potential compromise point that needs to have the reason for it being opened determined or the port needs to be closed. It was found that the SunRPC startup script had not been disabled. Since we cannot give a reason for having it enable the startup script has been added to the disabled. When configuring and performing an audit of a system it is useful to have a list of these files which are disabled or if they might be expected to be disabled but are left active a brief explanation of why is good. At the next audit this list can be used to determine if the configuration is correct. In the following you can now see that 21/tcp is no longer open.

```
root (www)>nmap 204.37.152.5
```

```
Starting nmap V. 2.54BETA28 ( www.insecure.org/nmap/ )  
Interesting ports on publicfw.giac.com (204.37.152.18):  
(The 1536 ports scanned but not shown below are in state: closed)
```

Port	State	Service
22/tcp	open	ssh
25/tcp	open	smtp
53/tcp	open	domain
70/tcp	open	gopher
80/tcp	open	http
111/tcp	open	sunrpc
113/tcp	open	auth
389/tcp	open	ldap
443/tcp	open	https
444/tcp	open	snpp
6000/tcp	open	X11
32771/tcp	open	sometimes-rpc5

```
Nmap run completed -- 1 IP address (1 host up) scanned in 6 seconds  
root (www)>
```

The log files are checked to see if the scans were detected and logged.

```
Aug 13 09:43:28 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56377 to 204.37.152.18 on unserved port 2028  
Aug 13 09:43:29 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 107  
Aug 13 09:43:30 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 639  
Aug 13 09:43:31 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 816  
Aug 13 09:43:36 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56377 to 204.37.152.18 on unserved port 816  
Aug 13 09:43:37 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 575  
Aug 13 09:43:38 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 2001  
Aug 13 09:43:39 publicfw.giac.com unix: securityalert: udp if=hme1  
from 204.37.152.20:56376 to 204.37.152.18 on unserved port 7009
```

The above entries in the log file indicate that the firewall detected the attempt to connect to various ports that had nothing listening on the port. This is what would be expected.

The following entry had a proxy process listening on the port but the firewall rules denied access to the proxy, since the telnet connection was attempted from the outside of the firewall this is also as expected.

```
Aug 13 09:01:56 publicfw.giac.com tn-gw[18241]: deny  
host=unknown/204.37.152.20 use of proxy
```

Design Modifications:

One modification to the design would be to provide a more centralized syslog function. With this centralized syslog it would be easier to correlate the events occurring on one subnet with those of the other subnets. The creation of a centralized syslog host provides this capability without having the task of moving the files from the individual syslog hosts to a main host and manually collating the files. Along with this the creation of ntp servers will be implemented to help with the correlation of logged events, all hosts will be using the same time for their logging. One approach would be to create a syslog host on the intranet protected by the intranet firewall from open attacks, however this would create a udp hole inbound through the intranet firewall which could provide a potential vulnerability for future exploitation. A decision would have to be made if the continuous creation of collated logs is worth the risk of this introduced hole or if the manual retrieval and merging of the logs is sufficient. An additional modification would be the addition of a VPN server with VPN client software. This addition would allow IPSec connections back to the GIAC Enterprises for GIAC employees on travel. The server receives VPN connections from the clients via any Internet connection. This would provide the employees on the road a way of staying connecting while on travel. We will also change the install for the firewalls to add the binding of the proxies to the appropriate interface.

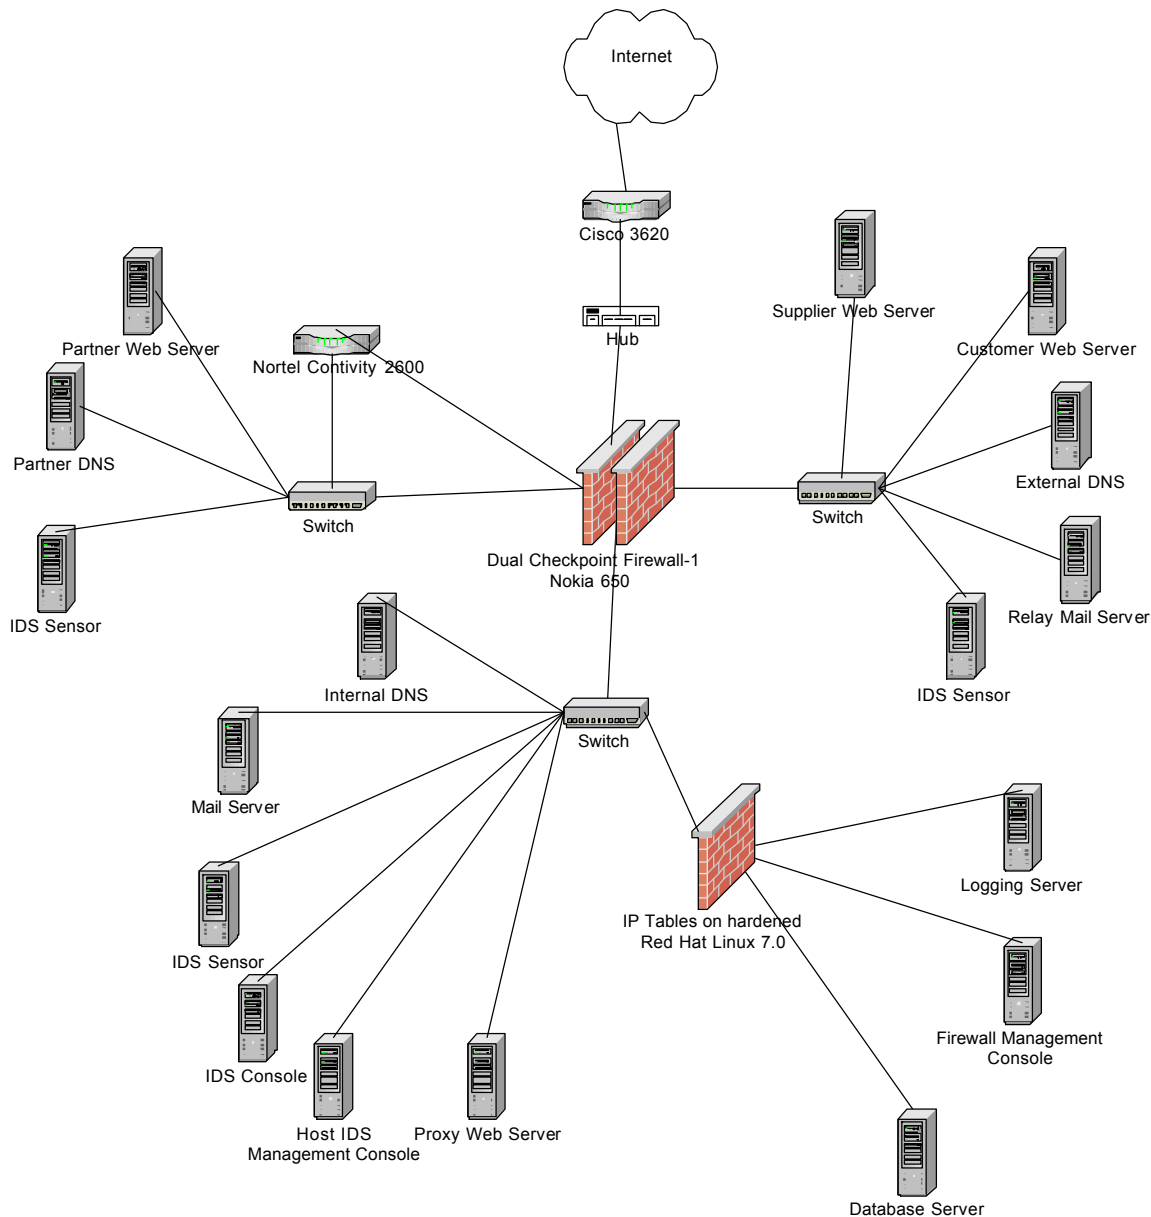
Question #4 – Attack Firewall using newly released security alert.

I have chosen the design by Matthew Brown,

http://www.sans.org/y2k/practical/Matthew_Brown_GCFW.zip,

for describing my attacks. The diagram below has been pasted from the design document done by Matthew Brown into this document.

© SANS Institute 2000 - 2005



Attack Against the firewall:

This design uses a Checkpoint Firewall 1 for the protection of the web servers accessed by the partners and suppliers. To find a vulnerability of the Checkpoint Firewall I performed a search at <http://icat.nist.gov/icat.cfm> for vulnerabilities for the past year. The ICAT site provided the following synopsis of a denial of service attack that the Checkpoint firewall is vulnerable to.

CAN-2001-0182

Summary:

Firewall-1 4.1 with a limited-IP license allows remote attackers to cause a denial of service by sending a large number of spoofed IP packets with various source addresses to the inside interface, which floods the console with warning messages and consumes CPU resources.

Published Before: 3/26/2001
Severity: Medium

In addition to the synopsis ICAT also provided this link <http://archives.neohapsis.com/archives/bugtrac/2001-01/0298.html> that provides a more detailed explanation of the vulnerability. Basically a denial of service can be achieved by sending spoofed address packets to the inside interface of the firewall. If the firewall has a limited number licenses, when the number of licenses is exhausted the firewall begins logging messages to the console. The messages logged include a list of all the IP addresses used in the license calculation. Eventually the list becomes so long that the list of IPs does not finish printing to the console before the next warning is issued and the firewall becomes so busy logging the license warnings that it can no longer process anything else and the firewall stops providing service and the denial of service is complete.

Denial of Service Attack:

The purpose of a Denial of Service (DOS) attack is to deny the client access to a particular resource. Examples of DOS are flooding a network with excessive traffic, disrupting connections between machines, or illegitimate uses of resources such as use of an anonymous ftp site for distributing unlicensed software. The methods for creating a DOS are the consumption of resources, destruction or alteration of configurations, or physical destruction of components [2]. I will concentrate on the first approach, consumption of scarce resources. Some of the scarce resources within a system are the network bandwidth, CPU processing power, disk space, and memory. The number of connections and user sessions is directly tied to the way these resources are managed within a system. SYN floods, UDP floods, and ICMP floods are directed at depleting these resources. Every connection made to a system uses a small amount of space that is reserved for the data structures it requires. A SYN Flood depletes these connection resources by creating TCP "half-open" connections. Sending a SYN with a spoofed IP creates a half-open connection. When the server responds with the required SYN-ACK back to the spoofed IP the third part of the three-way handshake never completes since the spoofed client will not respond to the SYN-ACK if it never sent the original SYN. At this point the server has created a data structure in memory for this connection and as more and more half-open connections are created the server runs out of space for creating new data structures and can no longer accept new connections [2]. A distributed denial of service (Ddos) attack is conducted by comprising many machines that are used as slaves or daemons by several master machines. The masters are also comprised machines and are controlled by the intruder performing the attack [3]. When the slave and master machines are compromised the TFN2K tools are installed. These tools provide the mechanism for launching various Ddos attacks utilizing SYN floods, UDP floods, ICMP floods and several others [4]. When these tools are inserted into the compromised machines the daemons communicate with the masters to let them know they are available and wait for commands from the masters. The SYN flood is the most difficult to guard against since it can be targeted at a legitimate port such as 80, http. It does have some difficulty since it requires that the demon attackers must be able to spoof their IP. Since the spoofed IP does not have to be of a machine that actually exists, we

can merely instruct the daemons to use the subnet and netmask they are on and generate the IPs to use for the spoofing so we can make the conclusion that there will be a sufficient number of hosts that fit into this category. So at the given time the intruder issues the commands to the masters and the masters in turn issue the commands to the slaves/daemons and the attack occurs.

Preventive measures for this type of attack are ingress and egress filtering on the routers along with modifying TCP parameter settings and patches for the specific OS being used. Most designs from knowledgeable security professionals will perform the ingress filtering for the unexpected subnets, your own IP ranges, the reserved IP ranges, and loopback along with egress filtering, only allowing source IPs from your subnet to exit your router. This helps prevent a daemon that has been installed on your subnets from sending out spoofed IP packets using the IP of other subnets. The TCP parameters in the OS that can be adjusted are the timeouts, the high water marks, and connection queue size. The default length for timeouts for the initialization and closing states of TCP connections is sometimes set to 4 or 5 minutes. This can be reduced to 1 minute so that the bogus connections will clear more quickly. The size of the connection queue can also be increased [5, 6].

Attack Against an Internal Machine:

Several factors must be considered when deciding on an internal machine to compromise. Some of these factors are the ease of access, the value of the machine, and the reason for doing the attack. This attack will be to gain access to a server that might yield a host for future snooping and attacks. The supplier web server may serve this purpose. To perform this attack on an internal machine there are two vulnerabilities available. One was found on the ICAT site <http://icat.nist.gov/icat.cfm> and requires fastmode to be configured on the firewall and the sending of malformed fragments to the firewall. Our hope is that since this firewall is passing traffic to a web server it was configured with fastmode to enhance web performance. The synopsis from ICAT follows:

CAN-2001-0082

Summary:

Check Point VPN-1/FireWall-1 4.1 SP2 with Fastmode enabled allows remote attackers to bypass access restrictions via malformed, fragmented packets.

Published Before: 2/12/2001

Severity: Medium

Additional details on this vulnerability and source code for exploiting it are available at <http://archives.neohapsis.com/archives/bugtrac/2000-01/0271.html>. The vulnerability occurs if fastmode is enabled and a TCP service is allowed access then all TCP services will be allowed through the firewall. So allowing access to a web server behind a firewall could allow access to other servers behind the firewall on other TCP ports such as telnet or ftp. I will use a connection to the telnet port to gain access to the server and deposit the programs for snooping passwords.

The second vulnerability uses port 259/UDP with a faked Reliable Data Protocol

header. This vulnerability was reported in and described at http://www.linuxsecurity.com/articles/firewalls_article-3312.html on 10 July 2001. This report stated, "The company that discovered the security hole, Inside Security GmbH said an attacker could add a faked RDP header to normal UDP traffic, allowing any content to be passed to port 259 on any remote host on either side of the firewall". The CERT reported, "that an intruder could use this vulnerability to launch certain kinds of denial of service attacks". This attack can be used as a distraction so the previous attack may be less likely to be noticed.

References

- 1) Chapman, D. Brent, and Elizabeth D. Zwicky. Building Internet Firewalls. O'Reilly & Associates, Inc. November 1995.
- 2) Denial of Service Attacks. The CERT Coordination Center Software Engineering Institute Carnegie Mellon University. Pittsburgh PA. June 4, 2001. http://www.cert.org/tech_tips/denial_of_service.html.
- 3) Results of the Distributed-Systems Intruder Tools Workshop. The CERT Coordination Center Software Engineering Institute Carnegie Mellon University. Pittsburgh PA. December 7, 1999. http://www.cert.org/reports/dsit_workshop-final.html.
- 4) CERT Advisory CA-1999-17 Denial-of Service Tools. The CERT Coordination Center Software Engineering Institute Carnegie Mellon University. Pittsburgh PA. March 3, 2000. <http://www.cert.org/advisories/CA-1999-17.html>.
- 5) Cockroft, Adrian, Richard Pettit. Sun Performance and Tuning. Sun Microsystems, Inc. Palo Alto, CA. 1998.
- 6) Solaris 2.x – Tuning Your TCP/IP Stack and More. <http://www.rvs.uni-hannover.de/people/voeckler/tune/EN/tune.html>.